

2026

Prediction, hype, and fear: Advancing Information Systems research in “the AI age”

Leslie Willcocks

Department of Management, London School of Economics, UK

Wendy Currie

Audencia Business School, Nantes, France

Daniel Schlagwein

Chulalongkorn University, Sasin School of Management, Thailand, The University of Sydney Business School, NSW, Australia

Jan Marco Leimeister

University of St.Gallen, Institute for Information Systems and Digital Business, St.Gallen, Switzerland, University of Kassel, Germany

Follow this and additional works at: <https://aisel.aisnet.org/jit>

Recommended Citation

Willcocks, Leslie; Currie, Wendy; Schlagwein, Daniel; and Leimeister, Jan Marco (2026) "Prediction, hype, and fear: Advancing Information Systems research in “the AI age”, *Journal of Information Technology*. Vol. 41: Iss. 2, Article 6.

DOI: 10.1177/02683962261460071

Available at: <https://aisel.aisnet.org/jit/vol41/iss2/6>

This material is brought to you by the AIS Affiliated Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in Journal of Information Technology by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Prediction, hype, and fear: Advancing Information Systems research in “the AI age”

Journal of Information Technology
2026, Vol. 41(2) 171–181
© Association for Information
Technology Trust 2026
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02683962261460071
journals.sagepub.com/jinf



Leslie Willcocks¹ , Wendy Currie² , Daniel Schlagwein^{3,4} and Jan Marco Leimeister^{5,6}

Abstract

Information Systems (IS) research is well-positioned but under-equipped to study technological futures at a time when claims about artificial intelligence (AI) are reshaping investment, policy, and public discourse. This perspective advances three arguments. First, IS scholarship should engage more systematically with digital futures, drawing on approaches for reasoning under uncertainty, such as Bayesian methods and established Futures Studies techniques, to distinguish prediction, projection, possibility, and hype. Second, technology hype is itself a legitimate object of IS research, and widely used frameworks such as the Gartner Hype Cycle appear limited in their ability to inform practice. Third, AI serves as a critical test case, combining heavy supply-side investment with unproven demand-side impact and unresolved questions of value and consequence. We propose four analytically distinct lenses for studying AI, namely, capability, adoption, value, and consequence, and identify two underexamined blind spots: bad actors deploying AI at scale and structural over-dependence on imperfect AI. We invite contributions to the Journal of Information Technology that examine how claims about technological futures are produced, circulated, institutionalized, resisted, and realized.

Keywords

digital futures, futures studies, technology hype, S-curve, Gartner Hype Cycle, artificial intelligence, IS research agenda

Introduction

Anticipating technological futures has always been a central, if unevenly developed, concern of Information Systems (IS) research (Schlagwein et al., 2025). This challenge has become even more acute in the current era, as claims about artificial intelligence (AI) rapidly reshape organizational investment, policy agendas, and public discourse. Recent advances in generative AI (genAI) have heightened expectations of transformative impact and uncertainty about the value and consequences of these systems. As these systems evolve from analytical tools toward systems capable of planning and executing multi-step tasks, they are increasingly embedded in organizational processes and decision-making.

Importantly, IS research has begun to recognize that such systems exhibit forms of agency that challenge long-standing assumptions about technology's role in organizational settings (Willcocks et al., 2026). Rather than functioning as passive tools, contemporary AI-enabled systems increasingly assume responsibility for task execution in ambiguous and uncertain conditions, shifting the human–system relationship from use to delegation (Baird

and Maruping, 2021). This shift has significant implications for how we conceptualize capability, control, and accountability in digitally mediated environments, particularly as AI systems assume more autonomous and consequential roles.

At the same time, there is growing evidence that these advances are accompanied by a widening gap between expectation and desired outcomes (Currie et al., 2026). Large-scale survey evidence indicates that although a substantial majority of organizations report using AI in at least one business function, most remain in the early stages

¹Department of Management, London School of Economics, UK

²Audencia Business School, Nantes, France

³Chulalongkorn University, Sasin School of Management, Thailand

⁴The University of Sydney Business School, NSW, Australia

⁵Institute for Information Systems and Digital Business, St.Gallen, Switzerland

⁶University of Kassel, Germany

Corresponding author:

Wendy Currie, Audencia Business School, 8 route de la Jonelière, BP31222, 44312 Nantes cedex 3, France.

Email: wcurrie@audencia.com

of deployment, with limited progress toward enterprise-level integration and measurable outcomes (Yee et al., 2025). For example, despite widespread adoption and increasing experimentation, few organizations report sustained enterprise-level performance gains, and many initiatives remain at the pilot stage. Accordingly, the disparity between AI's potential and its realized utility renders it pervasive yet still evolving, particularly given that "AI" is not a single construct but a diverse set of technologies with uneven maturity and impact (Chalmers et al., 2026).

These developments are further amplified by persistent dynamics of technological hype. The rise of genAI has been accompanied by rapidly escalating expectations, followed by growing recognition of technical, organizational, and governance constraints. Industry analyses suggest that despite significant investment, many organizations struggle to demonstrate consistent returns on AI initiatives and face ongoing challenges, including reliability, bias, and regulatory compliance. More broadly, the dynamics of technology hype, in which expectations expand faster than empirical validation, have long been recognized but remain only partially explained by dominant frameworks such as the Gartner Hype Cycle. Existing critiques suggest that such models oversimplify complex technological trajectories and offer limited guidance for understanding actual patterns of development and value realization (Dedehayir and Steinert, 2016). In the context of AI, hype is not merely discursive but actively shapes investment decisions, organizational behavior, and policy responses.

Taken together, these developments place IS research at a critical juncture. The field has traditionally excelled at identifying what works and what does not under stable conditions, often through empirical generalization and retrospective analysis. However, these approaches are less suited to environments marked by deep and rapidly evolving uncertainty, where both technological capabilities and the narratives surrounding them are changing simultaneously. In these settings, extrapolating from past evidence becomes increasingly unreliable, and distinctions among prediction, projection, possibility, and hype blur. As Spiegelhalter (2024) argues, uncertainty is not simply an objective property of the world but is shaped by assumptions, judgments, and incomplete knowledge, particularly in domains where future conditions cannot be directly observed.

This raises a fundamental question for the IS field: how should we study technological futures, particularly AI, when claims about capability, adoption, value, and consequences are advancing unevenly and often in tension with one another? In this perspective, the editors of this journal advance three connected arguments. First, IS scholarship should engage more systematically with digital futures, drawing on approaches to reasoning under uncertainty to better distinguish prediction, projection, possibility, and hype. Second, technology hype is a legitimate and

underexamined object of IS research, requiring more robust conceptual and empirical treatment than dominant frameworks currently provide. Third, AI, in all its manifestations, serves as a critical test case for examining these issues, combining rapid advances in technical capability with uneven adoption, ambiguous value, and significant concerns about governance and societal consequences. On this basis, we invite submissions to the Journal of Information Technology that address these challenges and advance our understanding of how technological futures are constructed, circulated, contested, and realized within organizational and societal contexts.

Technology futures beyond prediction

Humans are inherently predictive beings. They predict to anticipate, cope, survive, and gain some control over their environment and circumstances. Prediction means to say before, and in academic circles, it commonly implies inference from researched facts, principles, or accepted laws of nature. In IS, beyond the frequent use of statistics, probability, and "significance"-based studies, researchers have been wary of conducting studies that anticipate the future or establish guidelines that can be safely followed into the future. As studies, their findings can be projected into the future, but with ever-changing technologies and in dynamic contexts, this has very limited predictive value. The prevalence of these quantitative approaches in IS and their limitations have been well documented (e.g., Kucharski, 2025; Ziliak and McCloskey, 2008; Sen et al., 2021, and ensuing commentaries and response in JIT June 2022 issue).

Given the claims about Bayesian approaches for informing predictions under uncertainty and their widespread application across domains, including AI (Chivers, 2024; Poundstone, 2019), one might reasonably expect greater use of these methods in IS research, and we would encourage their broader deployment. In *The Art of Uncertainty*, Spiegelhalter (2024) highlights Bayes' theorem as a central mechanism for learning from experience, offering a coherent probabilistic framework for updating beliefs considering new evidence, particularly in contexts characterized by instability and uncertainty.

The IS field has had even less engagement with the field of Futures Studies. Partly, this has been due to an IS concern with establishing itself as a discipline. This has produced a preoccupation with Economics as a discipline it defers to for methodological and theoretical guidance. Partly, it will be a concern for the robustness of many of the methods cited in Futures Studies. Partly, it comes from an understandable reluctance to engage with the challenges and viability of attempting to predict and understand technological futures. However, this reluctance has become outdated.

First, the methods, modelling, and mathematization typical of Economics research have come under increasing

scrutiny regarding their usefulness and viability, not least from within the discipline. Senior Economists, John Kay and Mervyn King argue that, in the face of radical uncertainty, many research approaches in Economics have become increasingly outmoded and irrelevant (Kay and King, 2020). The authors criticize “bogus models” that rest on the fantasy of complete knowledge of the present and future (cp. chapter 20). As a prominent statistician, Spiegelhalter (2024) is also wary of maintaining practices in the face of increasing uncertainty and recommends more careful methods when predicting the future. For example, as we investigate futures, he suggests that assumptions about the structure and stability of models become more important; that uncertainty can be added to complex deterministic models by running ensembles from different initial states or with perturbed parameters; that insight can be gained by investigating a variety of prediction methods; but that humility is needed, particularly when the future is strongly influenced by human behavior or other unknowable factors. We suggest that such authors challenge IS researchers to re-examine their methods and offer useful guidelines on how to predict and carry out future research.

Second, Futures Studies has by now developed a large body of experience and knowledge, along with viable methods for studying the future. For example, over 25 years ago, Rescher (1998) provided a thorough handbook on conducting studies focused on the future and cited seven major predictive approaches in use. These range from the unsophisticated (judgmental estimation dependent on expert judgment) through formalized/inferential (trend projection; curve fitting; circumstantial analogy) to scientific/sophisticated (indicator coordination; law derivation, i.e., inference from accepted laws; phenomenological modelling using formal models). Even by 2003, the Technology Futures Analysis Methods Working Group could list 51 methods used in relevant published research (TFAM Working Group, 2003). This body of work has been extended since (Poll, 2024). There is a rich array of methods, know-how, and findings that IS researchers can draw upon.

Third, driven by the massive advances in and impacts of digital technologies, and by the extent to which their development and application, among many other drivers (see below), will shape our global future, “technology futures” has become one of the leading topics of our times (McRae, 2022; Skidelsky, 2023; Yee et al., 2025). In terms of technology futures, Willcocks et al. (2024) suggested that 10 “SMAC/BRAIDA”¹ technologies deserve ongoing IS research attention. These are social media, mobile, analytics, cloud computing, blockchain, robotics, automation of knowledge work (now widely called AI), internet of things, digital fabrication and modelling, and augmented reality. A watching brief is suggested for quantum computing, Web 3.0, etc., and bioengineering.

Using a comprehensive, methodical approach, Yee et al. (2025) plotted 13 technology trends by level of innovation,

public interest, equity investment, and adoption (as of mid-2024). AI scored highest on all four dimensions. Other emerging technologies expected to be influential include energy and sustainability technologies; mobility, digital trust and security, cloud and edge computing, bioengineering, robotics, immersive reality technologies, advanced connectivity, application-specific semiconductors, quantum technologies, space technologies, and emerging autonomous AI systems.

Fourth, with its strong academic underpinnings and distinctive knowledge base, IS as a field has an inside track on providing theorization and evidence-based, objective, and informed analysis and prescription. These are all too frequently lacking in practice, even in areas where the stakes are very high indeed, as is the case with the massive investments now being channelled into AI (Yee et al., 2025).

In short, IS studies are under-equipped to investigate technology futures, as they currently lack sufficient explicit methods, habits, training, and confidence for studying technology futures. At the same time, IS is well placed to carry out such research and can readily draw upon ample resources that IS researchers can shape to their own purposes and topics. The availability of data for such studies is unprecedented, and the urgency of the questions they address has never been greater. We encourage IS researchers to capitalize on this moment to invigorate the field. The opportunities outlined above offer fertile ground for those able to advance appropriate research methods and, with the right timing, secure access to valuable datasets for examining the evolution, applications, and impacts of these technologies. Yet, in doing so, IS researchers will quickly encounter the persistent technology hype that has accompanied predictions about digital futures since the earliest days of commercial computing.

Technology changes: Hype and fear

For over 50 years, a “hype-fear” lens has been a common way, especially in the media (subsequently enhanced by social media), in which computing, then automation, and successive waves of technological advance have been interpreted, through to today’s AI landscape, including generative models and increasingly autonomous (“agentic”) systems (Willcocks, 2020; Willcocks et al., 2026). Technology is portrayed as highly impactful, arriving quickly, and a force for good. Alternatively, this very power is seen as devastating for manual and skilled workforces, as well as for privacy, security, equality, control, ethics, ownership, surveillance, and even human existence itself (Currie et al., 2026; Deem and Warren, 2022; Skidelsky, 2023).

The 1995 Gartner Hype Cycle is a much-publicized, widely used framework that claims to address this technology hype. But does it? Gartner suggests a pattern over time of:

- (1) a technology trigger,
- (2) peak of inflated expectations,
- (3) a trough of disillusionment,
- (4) a slope of enlightenment, and
- (5) a plateau of productivity.

The Gartner Hype Cycle explains how a technology develops and is perceived over time. It uses seductive, arresting language, but how adequate is it as a model and explanation? In fact, though widely known, the “technology hero” journey it portrays has been heavily criticized. Looking at 20 years of technology, very few technologies move through this hype cycle, and most key technologies adopted since 2000 were not identified early in their adoption phases.

Mullany (2016) conducted a detailed analysis and reached several conclusions. A very large number of technology trends are “flashes in the pan,” and many technologies simply die. Gartner’s technical insight is often correct, but the implementation of the technology does not follow. A few core technical problems have been worked on for decades – for example, speech recognition, internet micropayments, and data analysis – yet they do not appear in hype cycle analyses. Some technologies keep receding into the future, while others make progress when no one is looking. Many major technical and socio-technical developments have flown under the hype cycle radar altogether, including x86 virtualization, NoSQL, Open Source, and Map/Reduce/Hadoop. The Economist (2024) endorsed the finding that the cycle is, in fact, a rarity. Per their analysis, only a fifth of technologies move through the cycle; many are adopted without transitioning through the steps (e.g., social media, cloud computing), and others fall by the wayside, with sixty percent that enter the “trough of disillusionment” not rising again. A high number of technologies are indeed flashes in the pan.

Dedehayir and Steinert (2016) provide a more academic underpinning for these conclusions. They have problems with the hype cycle’s theoretical derivation. First, the mathematical summation of an opinion-based, human-centric hype-expectations model and the classic technology S-curve model to form the singular hype cycle model seems questionable, as the underlying theories measure different phenomena.

Second, the Y-axis labelled “expectations” is imprecisely defined, allowing its meaning to shift. Indeed, when reviewing Gartner publications over time, we observe a shift between the concepts of expectations and visibility. However, there is no clear or robust framework for measuring these ideas. Furthermore, expectations reflect the perspectives of diverse stakeholder groups. This raises the question of whether expectations represent a composite view, and if so, how that view is quantified. A more nuanced analysis that accounts for the specific concerns of individual stakeholders is therefore essential.

In their empirical study of major technologies, Dedehayir and Steinert (2016) found that Gartner’s

expectation that a technology passes through the stages in 5 to 8 years did not hold. Far from the expected average “stage speed” (i.e., stage changes per year) of 0.625 to 1, the technologies they investigated had an average stage speed of only 0.23 stages per year, that is, an average hype cycle duration of 21.76 years, which exceeds even the longest prediction of Gartner’s long fuse technology cycles. Much of their empirical work calls into question the analytical capability of the hype cycle model. Overall, the work by these researchers provides a fine example of how IS scholars can engage with this subject area, and we recommend the proposals they make for progress in this area.

Our related analysis examines automation anxiety over the past 65 years (Currie et al., 2026). We, along with Dedehayir and Steinert (2016), value using a human-centric, hype-expectations model alongside the classic technology S-curve model. Looking at expectations, there is indeed a technology trigger. The first question is: In what ways are different stakeholders psychologically predisposed to react to or process the trigger? The next question is: What expectations are created about its growth, speed of deployment, and impacts, and which stakeholders expect to be positively or adversely affected? Part of these concerns how the technology triggers pathways and amplifies intended and unintended outcomes from different stakeholders, including the mainstream and new social media.

As the 21st century approached, the Y2K problem sparked a global three-year scare about a massive computer shutdown. The media and journalists fed off the (ancient) narrative of a potential millennial apocalypse; corporations grappled with public-relations narratives of readiness; and publishers issued several Y2K survival handbooks. In practice, very little happened on the target day. The key lesson regarding technology expectations is that perceived risk cannot be understood solely in terms of the likelihood of an event and its impact. Due to the powerful influence of communication systems, social media, and vested interests, the effects of technology and associated risks are increasingly shaped by the level of public attention. This aligns with Sandman’s (1993) concept of “outrage” – the emotional response that amplifies perception. IS researchers conducting similar work on technology and hype may well bear this in mind.

Then, we come to the applicability of the classic technology S-curve model. This suggests that a technology starts as costly, bulky, and little adopted. There follows a period when the technology’s fundamentals and its fit with organizations and problems are refined. Then, if this work is successful, rapid innovation and massive adoption may well follow, but eventually there will be a slowdown in significant improvement, and fewer new customers will be forthcoming. The peak of the S-curve declines, creating an opportunity for a better new technology with its own S-curve to come on top. Though not without criticism (e.g., timing of phases? Valid for all technologies?), the S-curve

forms a useful starting point for understanding how technologies roll out, including the current and future digital technologies.

Rogers (1962) expanded the curve to depict the diffusion of innovations within a culture over time. Diffusion is the process by which an innovation is communicated and adopted over time. Rogers' work is useful for underscoring that the innovation itself is not the only determinant of its "success." In the current attention given to AI, this means that people are mistaken if they believe AI's technological prowess is the only or even leading factor in whether it is adopted and institutionalized. Rogers' work suggests that three sets of factors shape the individual rate of adoption: perceived attributes of the innovation, communication channels, and the nature of the social system into which the technology is introduced. Within organizations, the rate of innovation is influenced by individual leader support for the change, six key characteristics of the organization, and the openness of the external system. The innovation process unfolds over time through five phases: knowledge, persuasion, decision, implementation, and confirmation. Augmenting this theory with more detailed, evidence-based theorization, Greenhalgh et al. (2004) and Willcocks et al. (2019) found very strong explanatory power when applied to 356 robotic process and cognitive automation cases followed during the 2014–19 period.

Both the classic S-curve and Rogers' theory of diffusion of innovation have limitations, but they provide a useful foundation for theorizing and for empirical work on emerging technologies. Additionally, Willcocks (2020) cautions toward a more realistic view of information and digital technology adoption, noting that these technologies are difficult to develop into workable systems and challenging to implement. Meanwhile, exploiting these technologies in organizational contexts has been consistently sobering. Thus, when looking at automation across decades, depending on the sector, only around 15–23% of corporations optimize business value, and even they accrue, on average, only 68% of the potential value. A common finding is that up to 65% of digital transformations seriously disappoint (Lamarre et al., 2023; Willcocks et al., 2024). Computer and digital technologies initially require adjustments and maintenance across the life cycle. Fitting them into an existing technical infrastructure is an understated and challenging task. Ensuring that the necessary managerial and technical skills are, in fact, present is not easily accomplished. As a result of such factors, the time horizons for developing a technology for full, optimal enterprise use can be quite long (Willcocks, 2020).

All this suggests that the hype-fear lens applied to AI tends to overlook major factors that will shape whether and how these technologies are adopted, and with what

consequences. We would encourage IS scholars to elaborate on these factors and take them much more into account in their research designs and studies.

The future of AI (and IS research on AI)

Let's take a closer look at the nexus between AI and the future and consider potential directions for IS research. Throughout the history of technology, a rich and enduring narrative has emerged, shaped by real events, that highlights technology's power, transformative potential, and dangers. In recent years, this perspective has evolved into a broader storyline that reflects uncertainty about a future increasingly shaped by digital technologies. A growing body of evidence, combined with vested interests and widespread anxiety, has lent weight to predictions of even more dramatic technological progress, along with significant consequences that may be both beneficial and harmful. After the Internet bubble that spanned from 1993 to 2001, these contrasting visions of utopia and dystopia gained momentum, leading to a more intense and widespread debate. For examples of this discourse, see the works of Kurzweil (2008), Carr (2011), and Keen (2015). Post-2015, these debates increasingly focused on "AI." In fact, we suggest that "AI" has become the poster child for all automation, and automation anxiety has crystallized into AI anxiety, a container for all digital "superagency" (Hoffman and Beato, 2025) as well as its perils and potential disasters (Currie et al., 2026).

Immediately, IS researchers, with their knowledge of digital technology, can help curb this language inflation, which has only served to expand the "AI" hype bubble. Language matters when we talk about these technologies. The term AI itself is highly misleading (Chalmers et al., 2026). Today, and for the foreseeable future, the relevant technology cannot be equated with the embodied cognition of human intelligence. More precisely, Polson and Scott (2018) call it, correctly, "a domain-specific illusion of intelligent behavior." What is referred to as AI today largely consists of data-driven machine learning and statistical inference techniques, often implemented through models such as neural networks, trained on large datasets and enabled by substantial computational power. As we and others have argued elsewhere, this, and what it produces, represents, at best, "weak AI." Despite suggestions to the contrary, we are very far from reaching the benchmark of "general human intelligence" (Bender and Hanna, 2025; Polson and Scott, 2018; Schlagwein and Willcocks, 2023).

Additionally, and unfortunately, the term AI has been all too often applied to and conflated with technologies that are not, under these definitions, AI, for example, robotic process automation, algorithms of all types, and many SMAC/BRAIDA technologies rebranded as AI. AI is widely used as a marketing term to promote companies, start-ups, and products that are often not, on any reasonable definition, AI-

based. Thus, one 2019 report found that two-fifths of Europe's artificial intelligence start-ups did not use any AI programs in their products (Ram, 2019). This set a trend for many subsequent misleading uses of the term "AI" (Bender and Hanna, 2025; Challapally et al., 2025).

The hype is further underscored when the impending "AI tsunami" is examined more closely. Based on emerging evidence, there has indeed been considerable investment in AI globally, especially since 2020. But these investments have mostly been on the supply side, for example, funding AI start-ups, developing more advanced AI, and expanding supportive capabilities such as cloud, storage, and large-scale data centers. For example, the four US-based "Big Tech" firms spent more than \$US 250 billion in 2025, mostly on AI infrastructure. Similar patterns of current and future investment can be seen in state-sponsored investments in China and Europe, as well as among private firms. But take-up by corporations has been much slower, and the return on investment (ROI) has been correspondingly low. Yee et al. (2025) reported that 78% of organizations were using AI in at least one business function, but while 92% of executives planned to invest more over the next 3 years, only 1% of leaders said their companies were fully mature in AI.

Challapally et al. (2025) found that despite \$US 30–40 billion in corporate investment in genAI, only 5% of organizations derived millions in value from integrated pilots, while the other 95% saw zero return. Over 80% of organizations had piloted or explored tools like ChatGPT and Copilot, with nearly 40% reporting deployment. However, these applications primarily enhance individual productivity rather than profit-and-loss (P&L) performance. Meanwhile, 60% of organizations evaluated enterprise-grade systems, 20% reached the pilot stage, and only 5% reached production. Most failed due to brittle workflows, a lack of contextual learning, and misalignment with day-to-day operations (Challapally et al., 2025).

Four lenses for IS research on AI: Capability, adoption, value, and consequence

A recurring move in the current AI debate is to collapse four distinct questions into a single claim. Public discussion routinely treats progress in technical capability as evidence of organizational adoption, adoption as evidence of realized value, and realized value as evidence of unstoppable societal consequences. We suggest that future IS research on AI is sharper when these are kept analytically distinct. We therefore identify four lenses for IS researchers studying AI and other emerging digital technologies. The capability lens asks what these systems can technically do, where they fail, and how their performance evolves. The adoption lens asks how they are taken up across individuals, teams, organizations, and sectors, and which factors accelerate, slow, or block diffusion. The value lens asks what organizational,

economic, and public value is actually realized, by whom, on what time horizon, and at what cost. The consequence lens asks what wider organizational, occupational, societal, geopolitical, and ecological effects follow, intended or otherwise.

The four lenses are not competing theories but ordered framings, and the choice should be driven by the research question rather than by the technology in fashion. We suggest the following heuristic. If the question concerns what a system can do or how performance scales, work in the capability lens and treat adoption and value as out of scope. If the question concerns who is using the technology, with what intensity, and why some organizations move faster than others, work in the adoption lens; here, Rogers' diffusion factors, Greenhalgh et al. (2004) organizational determinants, and absorptive capacity are central. If the question concerns realized business or public value, work in the value lens and resist the common slippage of citing pilot counts, license counts, or surveyed executive intent as evidence of value. If the question concerns labor, governance, ethics, geopolitics, or societal risk, work in the consequence lens, where AI imperfections, bad-actor threats, and over-dependence become first-order. Most strong studies will name a primary lens and use one or at most two adjacent lenses as scope conditions; weak studies tend to mix some or all four implicitly and discuss findings as if they applied to each.

But this cannot be the whole story about AI. Beyond insisting on more rigorous analysis of definitions and technologies, IS researchers have great opportunities to investigate the reasons for the markedly slow uptake of these technologies on the demand side, the challenges experienced, and why some organizations make faster progress and achieve positive returns. Part of this could involve investigating more deeply the many areas of concern about AI that have been raised. In our earlier JIT editorial, Schlagwein and Willcocks (2023) suggested a multi-point "AI imperfections" test that could be used to decide which attributes to research. Based on a review of extant AI research studies, the test identifies major concern areas: AI is brittle, opaque, greedy, shallow, hackable, amoral, biased, invasive, and fakeable. Beyond this guideline, there is already a large literature to draw upon regarding the ethical and social responsibility challenges, the technical challenges and limitations, and the wider political, work, and economic impacts (see, for example, Bender and Hanna, 2025; Carr, 2011; Deem and Warren, 2022; Smith, 2018; Vallor, 2024; Willcocks, 2024).

None of this obviates the need to advance understanding of how to realize the potential of these fast-developing, impressive technologies, as suggested by Hoffman and Beato (2025). However, if AI does deliver on its posited massive potential, in our view, two serious, underestimated blind spots require much more research than they are currently receiving.

The first blind spot, “bad actors with AI,” is the large and growing threat from bad (human) actors who, on current trajectories, will deploy AI and related advanced technologies at scale. The cybersecurity industry is large and growing fast for a reason. Yet much current AI scholarship treats the user as either a benign organization or an undifferentiated “society,” while the body of work on AI-assisted fraud, deepfake-driven social engineering, autonomous attack tooling, model theft, prompt injection, and the use of generative AI in influence operations remains thin in IS journals. We particularly welcome research on questions such as: how do organizations detect and respond to AI-augmented attacks; how do existing governance and audit regimes need to change when both attacker and defender are partly automated; how does the asymmetry between offensive and defensive uses of AI evolve; and what role can IS scholarship play in shaping cyber-norms, disclosure regimes, and procurement standards.

The second blind spot, “over-dependence on imperfect AI,” concerns the prospect that individuals, organizations, and entire societies will become structurally dependent on advanced digital systems that prove brittle, opaque, biased, etc., with adverse outcomes that surface only after the dependency is hard to reverse. This is the AI imperfections agenda taken seriously at scale. Productive questions include in which functions and sectors dependency is hardening fastest, and what the reversibility of those choices is; how organizations notice and recover from silent failures of automated decisions; how deskilling interacts with model drift over multi-year horizons; and what governance, redundancy, and human-in-the-loop arrangements actually preserve organizational and societal optionality. This blind spot is the natural complement to the consequence lens introduced above. IS is well placed to investigate it because it requires exactly the combination of technical understanding and organizational fieldwork that defines our field.

An unprecedented digital future?

Pursuing the theme of hyped technology, let us consider how IS researchers might deflate an even bigger techno-hype balloon. For present purposes, we will classify interested parties into three groups: those who are primarily proponents of AI, those who primarily register concerns, and those who see the advent, application, and futures of these technologies as requiring much more detailed and careful research. These are not exclusive categories, but we would see IS researchers as essentially in the third group.

The first two groups tend to subscribe to what we call “the unprecedented digital future motif,” reiterated, utilized, and magnified variously by journalists, social media users, hi-tech vendors, government reports, academics, and citizens alike. This motif is powerful partly because there is clear evidence of dramatic advances in technology, its capabilities, and impacts.

But have we not heard this before? If the technology is unprecedented, the storyline certainly is not. Media, social media, consultancies, and technology vendors have regularly hailed each emerging digital technology as distinctive, unprecedented in its impact, and a “breakthrough” that would replace previous technologies overnight. One can see this, for example, with the internet, mobile phones, social media, blockchain, cloud, more recently the metaverse, and, from 2020, genAI (e.g., ChatGPT), then agentic AI (e.g., AutoGPT). In practice, many technologies so heralded have not taken off. Others have taken much longer than anticipated to evolve in the face of deployment challenges and the need to fit with existing technologies, infrastructure, and business imperatives and processes. The result? If sometimes true, after years of reiteration, this storyline frequently sounds now all too much like marketing rather than reality.

To make this more concrete, Willcocks et al. (2019) suggest that, in the digital context, four domains need to be navigated: digital hype, digital capability, “useful” digital, and “strategic” digital. While AI as a broad field has long delivered value in areas such as statistical machine learning, the current wave of AI, particularly generative and foundation model approaches, remains in the experimental domain. As we have shown above, the 2025 evidence across many corporates suggests that few organizations have moved beyond the pilot stage, with limited evidence of sustained strategic application. In this sense, many of the technologies currently subsumed under the label “AI” have yet to demonstrate consistent business value at scale (Currie et al., 2026; Yee et al., 2025).

This “unprecedented digital future motif” has further contributed to a related narrative of a looming existential techno-apocalypse, strongly associated with the rise of genAI, particularly large-scale neural network and foundation model systems, including emerging agentic architectures. Scientists for Global Responsibility (2018) published results of a member survey in which over 80% perceived a medium to high chance of things going badly wrong with AI, while 96% wanted more AI regulation and 82% thought AI was more likely to create a dystopian rather than a utopian future. Such views entered the public imagination, with many influential voices, fearing the worst, expressing concerns about keeping AI socially responsible (see, for example, Harari, 2024; Kissinger et al., 2024; Tegmark, 2017). Thus, according to a Centre for AI Safety (2023) statement signed by hundreds of AI specialists: “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”

Whether framed in utopian or dystopian terms, such perspectives often overlook the extent to which continuity persists, especially in technology. As Edgerton (2008) aptly described, this is “the shock of the old.” In the case of information and digital technologies, this oversight is understandable. Organizations have invested heavily in

infrastructure, skills, and applications that are deeply embedded in their operations. These technologies not only create dependencies but also generate significant value. However, integrating new systems into existing technical and organizational frameworks is time-consuming and resource-intensive. In the face of even more change, organizations may, and do, run out of absorption capacity (Willcocks et al., 2024). Even more challenging is the regulation of emerging technologies. Evidence from financial markets shows that regulators with limited understanding of high-frequency trading (HFT) tend to shift from decision optimization to satisficing behavior when overseeing such systems (Currie et al., 2022; Currie and Seddon, 2021).

Notwithstanding these challenges, it is important to recognize that digital technologies may not be the sole, or even the primary, drivers of the future. A wider perspective reveals a constellation of influential forces, including demographic change, climate pressures, resource scarcity (such as energy and water), environmental degradation, geopolitical tensions, evolving governance models, rising nationalism, armed conflict, migration patterns, and shifts in global trade dynamics (Boston Consulting Group, 2025; Guillen, 2020; McRae, 2022; Tai, H., 2025).

Returning to our main argument, Atkinson and Moschella (2024) offer an alternative view that, apart from deepfakes, most AI fears are still speculative; many seem all too familiar, often inherited from past scares, and many (they list 12 major fears) seem manageable. They also point out that the 1900–1960 period experienced technological and other innovations across industry, work, and society that had much more profound impacts than those of the following years. They also show that the pace of technological change is not accelerating and suggest that overestimating the speed leads to exaggerations, fears, anxieties, and an anti-innovation mindset. We would summarize this by saying that today's concerns all too often display an ahistorical, recency bias that contributes to much additional, unnecessary anxiety about the future (Currie et al., 2026).

Conclusion

IS researchers, drawing on their methodological expertise, technological and digital competencies, and accumulated theoretical and empirical knowledge, are well-positioned to challenge prevailing technology hype. They can produce rigorous, evidence-based analyses to critically assess and ultimately expose AI-related hype, grounded in systematic, reality-based research.

Consider technologies such as genAI (e.g., ChatGPT), which are often framed as transformative breakthroughs expected to shape an unprecedented digital future. Yet, our present situation resembles many earlier moments in the history of technological innovation. As with previous

technological waves, substantial excitement and considerable anxiety have accumulated around the advances attributed to genAI. It is widely agreed that its speed and the range of tasks it can accomplish are highly impressive. Many are already experimenting with this technology in its early stages. While the underlying development costs are substantial, these technologies are inexpensive at the point of use, widely accessible, and are being tested for usability, task effectiveness, potential to address major challenges, and reliability, consistency, and transformative impact, rather than simply being another practical tool. However, concerns have emerged across at least six key areas (Currie et al., 2026): Will it displace jobs or tasks? Will it cause widespread deskilling? Could it lead to increased but riskier outsourcing? Will its drawbacks outweigh its benefits? Is it contributing to technostress? What are the risks to data and information security? How real are the geopolitical, societal, indeed existential threats? For now, these questions, and doubtless many others, remain unresolved. But they present prime opportunities for IS researchers to utilize their distinctive experiences, knowledge bases, and research skills.

As editors, we close by directly stating the kinds of papers JIT particularly seeks to attract under the broader “digital futures” agenda set out in Schlagwein et al. (2025), which this perspective extends. We welcome IS studies that explicitly address uncertainty in technology trajectories, including (but not limited to) Bayesian approaches that update beliefs under genuine, non-stationary uncertainty rather than relying on null-hypothesis significance testing; IS engagements with Futures Studies that apply, rather than merely cite and conceptually argue the benefits of, the methods catalogued by Rescher (1998), the TFAM Working Group (2003), and Poll (2024); empirical critiques of techno-hype frameworks, especially the Gartner Hype Cycle, together with constructive alternatives that decouple the human-expectations curve from the S-curve and integrate diffusion theory and Sandman’s “outrage”; and studies that use the four-lens framework (capability, adoption, value, and consequence) to keep AI claims analytically separate and to reality-test current AI ROI narratives against multi-year, multi-organization evidence.

We are equally interested in research on the two AI blind spots flagged above, namely, bad actors deploying AI at scale and structural overreliance on imperfect AI – including questions of governance, audit and reversibility; in comparative work on how AI hype and adoption vary across global contexts, particularly between Big Tech-dominated regions and the Global South, and on how geopolitical and industrial-policy framings shape what counts as a “digital future”; and in work that treats hype itself as an IS object of study, asking how technology expectations are produced, circulated, institutionalized, resisted, and realized, rather than papers that either amplify or merely debunk hype. Across these directions, we particularly welcome

submissions that integrate technological understanding with organizational, societal, and methodological rigor, thereby enabling readers to distinguish plausible futures from projections, promises, and fears. In part, this is the role that JIT exists to fulfill, and it remains central to the journal's innovative mission.

ORCID iDs

Leslie Willcocks  <https://orcid.org/0000-0003-2572-9554>

Wendy Currie  <https://orcid.org/0000-0001-9198-6933>

Daniel Schlagwein  <https://orcid.org/0000-0002-1591-4660>

Jan Marco Leimeister  <https://orcid.org/0000-0002-1990-2894>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Note

1. Social, Mobile, Analytics, and Cloud (SMAC) Bridging Responsible AI Divides and Adoption (BRAIDA).

References

- Atkinson R and Moschella D (2024) *Technology, Fears and Scapegoats: 40 Myths about Jobs, AI, and Today's Innovation Economy*. Palgrave Macmillan.
- Baird A and Maruping LM (2021) The next generation of research on IS use: a theoretical framework of delegation to and from agentic IS artifacts. *MIS Quarterly* 45(1): 315–341. <https://doi.org/10.25300/misq/2021/15882>
- Bender E and Hanna A (2025) *The AI con: How to fight Big Tech's Hype and Create the Future We Want*. The Bodley Head.
- Boston Consulting Group (2025) *Ten Forces Reshaping Global Business*. BCG.
- Carr N (2011) *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*. Atlantic Books.
- Centre for AI Safety (2023) Statement on AI risk. <https://www.safe.ai/statement-on-ai-risk> and <https://safe.ai/work/impact-report/2023>
- Challapally A, Pease C, Raskar R, et al. (2025) *State of AI in Business: The Gen AI Divide*. MIT NANDA.
- Chalmers D, Hunt R, Pachidi S, et al. (2026) The acceleration of artificial intelligence: rethinking organization and work in an era of rapid technological change. *J. Manage. Stud* 63: 285–314. <https://doi.org/10.1111/joms.70063>
- Chivers T (2024) *Everything Is Predictable: How Bayes' Remarkable Theorem Explains the World*. Weidenfeld & Nicolson.
- Currie W and Seddon JJM (2021) Exploring technological instantiation of regulatory practices in entangled financial markets. *Journal of Information Technology* 37(1): 31–50. <https://doi.org/10.1177/02683962211027308>
- Currie WL, Seddon JJM and Van Vliet B (2022) From decision optimization to satisficing: regulation of automated trading in the US financial markets. *Information and Management* 59(8): 103721. <https://doi.org/10.1016/j.im.2022.103721>
- Currie WL, Seddon JJM and Willcocks LP (2026) Beyond job loss: mapping automation anxiety across five eras of computing. In: IADIS Conference, Zagreb, Croatia, 7-9 March 2026. Awarded Best Paper.
- Dedehayir O and Steinert M (2016) The hype cycle model: a review and future directions. *Technological Forecasting and Social Change* 108: 28–41. <https://doi.org/10.1016/j.techfore.2016.04.005>
- Deem M and Warren P (2022) *AI on Trial*. Bloomsbury Professional.
- Edgerton D (2008) *The Shock of the Old: Technology and Global History since 1900*. Profile Books.
- Greenhalgh T, Robert G, MacFarlane F, et al. (2004) Diffusion of innovation in service organizations: systematic review and recommendations. *The Milbank Quarterly* 82(4): 581–629. <https://doi.org/10.1111/j.0887-378X.2004.00325.x>
- Guillen M (2020) *2030: How Today's Biggest Trends Will Collide and Reshape the Future of Everything*. Flint.
- Harari Y (2024) *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Fern Press.
- Hoffman R and Beato G (2025) *Superagency: What Could Possibly Go Right With Our AI Future*. Authors Equity.
- Kay J and King M (2020) *Radical Uncertainty: Decision-Making for an Unknowable Future*. Bridge Street Press.
- Keen A (2015) *The Internet Is Not the Answer*. Atlantic Monthly Press.
- Kissinger H, Schmidt E and Huttenlocher D (2024) *The Age of AI*. Little, Brown and Company.
- Kucharski A (2025) *Proof: The Uncertain Science of Certainty*. Profile Books.
- Kurzweil R (2008) *The Singularity Is Near: When Humans Transcend Biology*. Penguin Books.
- Lamarre E, Smaje K and Zimmel R (2023) *Rewired*. Wiley.
- McRae H (2022) *The World in 2050: How to Think about the Future*. Bloomsbury.
- Mullany M (2016) Eight lessons from 20 years of the Gartner hype cycle. *LinkedIn*. <https://www.linkedin.com/pulse/8-lessons-from-20-years-hype-cycles-michael-mullany/>
- Poll R (ed) (2024) *Handbook of Futures Studies*. Edward Elgar Publishing.
- Polson N and Scott J (2018) *AIQ: How Artificial Intelligence Works and How We Can Harness its Power for a Better World*. Bantam Press.
- Poundstone W (2019) *How to Predict Everything: the Formula Transforming What We Know About Life and the Universe*. Oneworld.
- Ram E (2019). *Europe's AI Start-Ups Often Do Not Use AI, Study Finds*. Financial Times.
- Rescher N (1998) *Predicting the Future: An Introduction to the Theory of Forecasting*. State University of New York Press.

- Rogers E (1962) *Diffusion of Innovations*. 5th edition. Free Press.
- Sandman PM (1993) *Responding to Community Outrage: Strategies for Effective Risk Communication*. American Industrial Hygiene Association.
- Schlagwein D and Willcocks L (2023) ChatGPT et al.: The ethics of using (generative) artificial intelligence in research and science. *Journal of Information Technology* 38(3): 232–238. <https://doi.org/10.1177/02683962231200411>
- Schlagwein D, Currie W, Leimeister J, et al. (2025) Digital futures: definition (what), importance (why) and methods (how). *Journal of Information Technology* 40(1): 2–8. <https://doi.org/10.1177/02683962241301544>
- Scientists for Social Responsibility (2018) Scientists' AI poll points to a dystopian future, less control and a high chance of mistakes [Press release]. Available at: <https://www.sgr.org.uk/resources/scientists-ai-poll-points-dystopian-future-less-control-high-chance-mistakes>.
- Sen A, Smith G and Van Note C (2021) Statistical significance versus practical importance in information systems research. *Journal of Information Technology* 37(4): 286–302. <https://doi.org/10.1177/02683962211062236>
- Skidelsky R (2023) *The Machine Age: An Idea, a History, a Warning*. Allen Lane.
- Smith G (2018) *The AI Delusion*. Oxford University Press.
- Spiegelhalter D (2024) *The Art of Uncertainty: How to Navigate Chance, Ignorance, Risk and Luck*. Pelican.
- Tai H (2025) What forces will shape the next 60 years? *McKinsey Quarterly* 3: 68–69.
- Technology Futures Analysis Methods Working Group (TFAM) (2003) Technology futures analysis: toward integration of the field and new methods. *Technological Forecasting and Social Change* 71: 287–303.
- Tegmark M (2017) *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf.
- The Economist. (2024) *Artificial Intelligence is Losing Hype*. The Economist.
- Vallor S (2024) *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking*. Oxford University Press.
- Willcocks L (2020) Robo-apocalypse cancelled? Reframing the automation and future of work debate. *Journal of Information Technology* 35(4): 286–302. <https://doi.org/10.1177/0268396220925830>
- Willcocks L (2024) Automation, digitalization and the future of work: a critical review. *Journal of Electronic Business and Digital Economics* 3(2): 184–199. <https://doi.org/10.1108/jebde-09-2023-0018>
- Willcocks L, Hindle J and Lacity M (2019) *Becoming Strategic with Robotic Process Automation*. SB Publishing.
- Willcocks L, Hindle J, Stanton M, et al. (2024) *Maximizing Value with Automation and Digital Transformation*. Palgrave Macmillan.
- Willcocks L, Currie WL and Seddon JJM (2026) Deconstructing the 'AI meme': beyond hype, fear, and digital transformation. In: IADIS Conference, Zagreb, Croatia, 7-9 March 2026.
- Yee L, Chui M, Roberts R, et al. (2025) *Technology Trends Outlook 2025*. McKinsey & Company.
- Ziliak ST and McCloskey DN (2008) *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives*. University of Michigan Press.

Author biographies

Leslie Willcocks is Professor (emeritus) in the Department of Management at the London School of Economics and Political Science. Leslie has a global reputation for his work in robotic process automation, AI, cognitive automation and the future of work, digital innovation, outsourcing, global management strategy, organizational change, IT management, and managing digital business. He is co-author of 75 books on these subjects, and has published over 240 refereed papers while his work appears in major media outlets such as Forbes magazine, HBR Online. He keynotes regularly at international conferences, has delivered executive programmes globally for some 30 years and has been retained as adviser and expert witness by major corporations and government institutions. Recent and forthcoming publications include: *Transformation in Global Outsourcing: The Sourcing of Digital Technologies and Services* (2024); *Global Business Strategy in Context* (2021); *Maximizing Value With Automation and Digital Transformation: A Realist's Guide* (2024). Website: www.roboticandcognitiveautomation.com

Wendy Currie is a Full Professor of Information Systems and Management at Audencia Business School, France. She is the joint Editor-in-Chief of the *Journal of Information Technology*. Her research focuses on the intersection of policy and information technology, with particular attention to emerging technologies, digital transformation, and disruptive automation in sectors such as finance and health-care. She also conducts comparative analyses of technology performance and outcomes across countries. Much of her work draws on institutional theory and stakeholder analysis. She has over 100 peer-reviewed papers in IS, management, and public administration outlets. Before her current role, she launched the MSc in Information Systems and Management at Warwick Business School and was Head of Department. She also served as the founding Editor-in-Chief of Health Policy and Technology (HPT), during which the journal achieved its first impact factor, marking a significant milestone in its academic visibility and influence. Her research has received funding from the European Commission, the Engineering and Physical Sciences Research Council, and the Economic and Social Research Council.

Daniel Schlagwein is Professor at the Chulalongkorn University Sasin School of Management and the University of Sydney Business School, is joint Editor-in-Chief of the *Journal of Information Technology*, President of the AIS

SIG Grounded Theory Methodology (SIG GTM), and an internationally leading researcher on digital work and digital nomadism. His research has resulted in over 90 peer-reviewed papers (e.g. *Information Systems Journal*, *Information Systems Research*, *Journal of Business Ethics*, *Journal of the Association for Information Systems*, *Technology and Society*) and has been featured in the media (e.g. ABC, The Age, El País, Sydney Morning Herald, Wired), has been the basis for advice to industry and governments (e.g. on the Thai “digital nomad visa” 2024). His work has received “Best Paper of the Year” awards from the Association for Information Systems (AIS) and the Association for Information Science and Technology (ASIS&T), and he has been a member of the AIS College of Senior Scholars (CSS) since 2020.

Jan Marco Leimeister (JanMarco.Leimeister@unisg.ch) is Chair Professor and Managing Director of the Institute of Information Systems and Digital Business at the University of St. Gallen, Switzerland. He is also Director at the

Research Center for IS Design at the University of Kassel, Germany, and an AIS Fellow. His work covers digital business, AI-driven value creation, digital service management, crowdsourcing, digital Work, digital learning services, and collaboration engineering. He has been internationally recognized for outstanding research, teaching, and education. He ranks repeatedly since 2009 among the top 1% of the most productive researchers and professors in Business Administration in the German-speaking area. He has served on the AIS Leadership Council, is Co-Editor in Chief of the *Journal of Information Technology*, and serves on the Editorial Board of the *Journal of Management Information Systems* and *Information Systems Research*. Dr. Leimeister is also an entrepreneur, and acts as senior advisor, board member, and keynote speaker for national and international organizations. His research has been funded by German, Swiss, and European Union ministries as well as by industry, foundations, and NGOs.