AI for Knowledge Creation, Curation, and Consumption in Context

David Schwartz
Bar-Ilan University, Israel, david.schwartz@biu.ac.il

Dov Te'eni
Tel Aviv University, teeni@post.tau.ac.il

AI for Knowledge Creation, Curation, and Consumption in Context

David Schwartz,[1] **Dov Te'eni**[2]

[1]Bar-Ilan University, Israel, david.schwartz@biu.ac.il
[2]Tel Aviv University, Israel, teeni@tau.ac.il

1 Introduction

Will this be the last editorial that we write or need to write? Based on the opinions presented in this special issue, it is plausible that AI could fully overtake the journal publication process, beginning with AI-generated reviews, as discussed by some of our contributors, until a fully AI-driven journal is realized, including editorials. However, other views that you are about to encounter in this special issue would suggest otherwise—that our role as editors will remain essential, with a wide range of justifications raised indicating the importance of maintaining high levels of human involvement in all aspects of the production, assessment, and publication of academic journal articles. But in questioning our roles in the journal ecosystem, we beg the fundamental question of what form future knowledge production *should* take, a fundamental issue taken up by a number of our contributors. Other opinions that you are about to read examine the current roles of author, editor, reviewer, publisher, and the eventual reader as supported by existing institutional dynamics that may themselves give way to new epistemological foundations and units of knowledge produced, validated, and shared in novel ways. The winds of change are blowing, and we had better start to figure out how to set our sails.

Welcome to this special issue of the *Journal of the Association for Information Systems* (JAIS) commemorating the life and work of JAIS founding editor Phillip Ein-Dor. As we introduce the fascinating topics that have been addressed by some of our thought leaders, we do so with the underlying subtext: "What would Phillip have thought?" We'll get back to that toward the end, but let's begin with an overview of

how academic journals operate, why they work the way they do, and what challenges lie ahead.

AI will inevitably impact academic journals. Academic journals can be seen as systems designed to publish academic articles that impact research and practice. Journals source, legitimate, and curate knowledge for consumption by researchers and practitioners, functioning in a particular context. These systems include interrelated editorial, production, and distribution processes operating in a dynamic ecosystem in which research and practice are continually evolving. Importantly, these systems can adapt and learn from feedback. The recent advancements in AI have the potential to strategically and operationally impact the structure and processes of academic journals, such as their economic models, social and ethical procedures, learning capabilities, knowledge dissemination and sharing mechanisms, and intellectual contributions.

There is, however, much uncertainty about whether and how to seize these new opportunities considering the risks involved. The confusion and urgency to act experienced by academic journals is evident in the very diverse and often conflicting policies that journals are now issuing. Some journals are rapidly moving to announce restrictive guidelines and ethical standards regarding the use of large language models (LLM) and generative pretrained transformers (GPT) in the *writing* process. Other journals are more liberal, not even requiring reports of using AI, provided it was not used for modeling. Yet other journals prohibit the use of any generative AI (gAI) in the review process. The Committee on Publication Ethics (COPE), released a clear position statement covering two main points: (1) AI cannot be assigned authorship and (2) authors who

use AI tools in the writing of a manuscript, must disclose in the paper's materials and methods section (or similar section) how the AI tool was used and which tool was used (COPE, 2023). Mainstream research publications *Science* (Thorp, 2023) and *Nature* (Nature, 2023; Stokel-Walker, 2023) have taken similar positions. The AIS Code for Research Conduct (AIS, 2014), last updated in 2014, is clearly in need of an update in this regard. Both gAI and more established machine learning (ML) models have the potential to drastically alter many aspects of academic scholarship—from knowledge production to assessment, verification, dissemination, and beyond. The need for clarification and analysis of these new opportunities and risks for our journals is obvious and urgent.

In this special issue (SI) of opinion pieces on the role and impact of AI, we go beyond the question of AI authorship (Dwivedi et al., 2023) and reach into the heart of the journal production process and value chain, examining aspects ranging from the use of AI by reviewers (Checco et al., 2021) to editorial decisions based on AI (Yuan et al., 2022) to fundamental changes in the journal ecosystem in a world in which many core scholarly activities will increasingly be conducted using gAI (Susarla et al., 2023). We begin our editorial with the process of compiling this SI and then present an overview of the individual contributions and offer a consolidated thematic view that integrates the contributions.

2 The Challenge and Response

For this SI, we approached editors of information systems (IS) journals, including all JAIS editors-in-chief (EICs) that succeeded founding EIC Phillip Ein-Dor. A few of the invited authors requested co-authorship with colleagues who were already collaborating on related work. The authors were asked to consider several issues that academic journals will potentially face in the coming years focused around the role of AI in the *creation, curation, and consumption of knowledge*, for which academic research journals are a fundamental building block of science, innovation, and progress. Journal articles have historically created a "conversation" between generations of researchers and practitioners. (Eva et al., 2023; Steingard & Linacre, 2023). Manuscripts reporting new research are based on prior literature, taking care to cite research considered to be thorough, trustworthy, and of rigorous high standards—but that fundamental knowledge creation process is itself poised to change.

The contributed opinion pieces fall into several categories: the creation of new research and integration of extant research; curation through the editorial process, notably reviewing; consumption, including

the dissemination of curated materials; and context, including the scientific domain, notably IS and the journals' environment and infrastructure. Figure 1 illustrates the processes of a journal ecosystem, including the management of journals, which was addressed by several authors. The editorial process is at the core of curation in academic journal publication. The definition of journal goals, relevant topics, and acceptable methodologies, as well as the editorial board composition and the roles of senior and associate editors and reviewers, are all shaped by editorial leadership, resources, and infrastructure. (Shmueli et al., 2023, Eva et al., 2023, Steingard & Linacre, 2023, Yuan et al., 2022).

A number of these central curation issues are discussed in "Reimagining the Journal Editorial Process: An AI-Augmented versus an AI-Driven Future" by Galit Shmueli and Soumya Ray, with issues of responsible use at all levels addressed by Shirley Gregor in "Responsible Artificial Intelligence and Journal Publishing." These papers also examine the dissemination and *consumption* of knowledge, as well as *journal management*, which are tightly linked to the editorial process.

Within the creation process, we singled out reviewing, which attracted much of the authors' attention. *Reviewers and the review process* are the engine that drives manuscripts from submission through evaluation, critique, improvement, and recommendation, and ultimately provide editorial leadership with the assessment and information needed to make a publication decision. Human reviewers are also the primary bottleneck in the publication process. Identifying appropriate reviewers, soliciting and receiving review commitments, and receiving thorough reviews in a timely manner are all challenges that editorial teams face on a daily basis (Checco et al., 2021, Kumar et al., 2023, Pradhan et al., 2021, Shmueli et al., 2023). Many of these issues are discussed in "Peer Review in the Age of Generative AI" by Atreyi Kankanhalli; "The Other Reviewer: RoboReviewer" by Ron Weber; and "Human-in-the-Loop AI Reviewing: Feasibility, Opportunities and Risks" by Iddo Drori and Dov Te'eni.

Focusing on theorizing and literature reviews, discussions of the earliest stages of the knowledge creation process itself—characterized by knowledge discovery and the integration of extant knowledge, analysis of the findings, and the writing of manuscripts—are found in Sirkka Jarvenpaa and Stefan Klein's "New Frontiers in Information Systems Theorizing: Human-gAI Collaboration" and Ojelanki Ngwenyama and Frantz Rowe's "Should We Collaborate with AI to Elaborate Research Articles and Literature Reviews?" These two papers each provide important insights into what lies ahead.
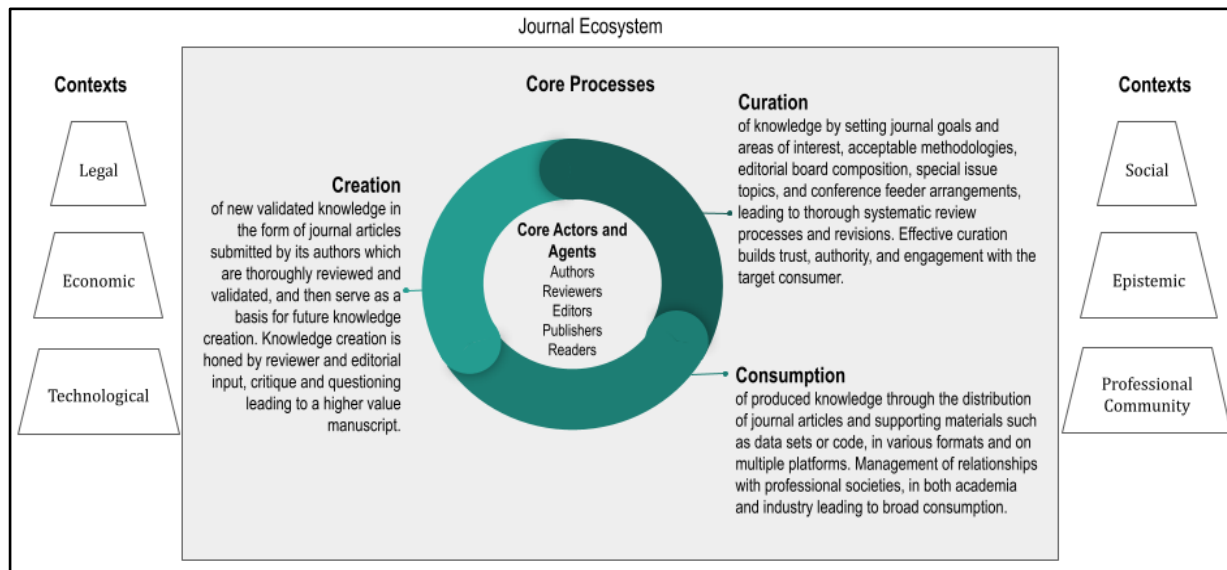
Figure 1. Academic Journal Ecosystems Drive Creation, Curation, and Consumption of Knowledge While Considering Multiple Contexts

A clearly identifiable group of papers in this SI touch on new ways of packaging and consuming journal knowledge and the broad *context* in which journals exist. These contributions cover the impact of AI-augmented journal processes on the consumption of knowledge and the epistemic and economic infrastructure. They include "Epistemic Infrastructure: The Role of Scientific Journals in the Age of Generative AI" by Youngjin Yoo; "Causal Knowledge Analytics: Knowledge Maps" by Richard Watson, Yuanyuan Song, Xia Zhao, and Jane Webster on the knowledge structure needed for smart editorial processes; "Toward Democratizing Knowledge through Academic Journals: A Proposal for Harnessing the Power of AI and Human Collaboration" [1] by Suprateek Sarker, Anjana Susarla, Ram Gopal and Jason Thatcher; and "Digital Transformation of Academic Publishing: A Call for the Decentralization and Democratization of Academic Journals" by Michel Avital.

3   A Consolidated Thematic View

Adding AI to the creation-curation-consumption-context elements composing the journal ecosystem (Figure 1) highlights four thematic challenges addressed by the papers in this SI: (1) the delegation of responsibilities and control, (2) interpretability, (3) accountability and responsible AI, and (4) impact. Following Ågerfalk (2020), we take an IS perspective to make sense of AI's roles and impact, looking at each element of the journal ecosystem as well as the entire system. Ågerfalk argues that the impact of AI can be analyzed by viewing it as a system that stands out in terms of its reliance on ML, its extensive datafication, and its complex and interrelated infrastructural resources. It is assumed that the human will not only remain in the loop but will also retain control of the system and delegate responsibilities. The IS perspective draws attention to the messages exchanged between subsystems in the journal ecosystem, including interpretable messages exchanged between the human and the machine. Notably, messages are socially meaningful units that are received and understood in a particular context, and they have consequences that affect people—consequences for which the human agents controlling the system must be held accountable (Ågerfalk, 2020). The *delegation* of responsibilities from the human to the machine depends and is dependent on the issues of *interpretability, impact* (intended and unintended consequences), and *accountability*. We review each of these four issues in light of the 11 opinion papers in the SI and formulate them as design dilemmas for future designs of AI-augmented journal systems.

3.1   The Delegation of Responsibilities and Control

The first dilemma journals will have to resolve is the mode of human-AI collaboration they wish to implement. All the papers in the SI assume or explicitly state that human agents will be kept in the loop and that they will

---

[1] A few of the opinions could be placed in multiple categories, e.g., Sarker et al. concentrated on reviewing and

general impact, but we decided to choose only one fitting category.

maintain control and be held accountable for the products and impacts of the AI-augmented journal ecosystem. The authors differ in their opinions regarding the subsystems (tasks) that can be delegated to the machine to be fully or partially automated, and, moreover, the criteria for delegating responsibilities to either the human or the machine. The question of delegation is not new to designers of human-computer interaction, but it has become more challenging. New AI technologies, such as gAI, present a greater range of capabilities with greater uncertainty and complexity, which makes it difficult to assess the machine's capabilities. The greater range, uncertainty, and complexity of machine capabilities also make it harder to assess the risks of negative social impacts, making it necessary to introduce new criteria for delegation. The complexity and uncertainty of the greater capabilities and greater risks require a new examination of how and when to delegate without surrendering human control and accountability.

Alternative delegation decisions dictate alternative human-machine configurations, which, in turn, pose distinct opportunities and risks. Kankanhalli examines the different evaluation tasks in reviewing papers, i.e., evaluating different aspects of a paper, on the basis of relative performance. She finds that for tasks such as judging novelty, significance, and plagiarism, AI tools perform worse than humans. Additional considerations that may be considered in delegation include the introduction and intensification of biases. Furthermore, delegation should be subject to ethical considerations, including the need for explainability and accountability. Weber, too, looks at delegation in reviewing papers. He considers the opportunities in automating (rather than augmenting) the generation of a review and proceeds to discuss the delegation of tasks between humans in the roles of editors, authors, and reviewers and the machine (RoboReviewer) by asking what competencies are needed to perform these roles and how machines can provide these competencies. Drori and Te'eni give us a microlevel view of how a fully delegated LLM-based review process compares to a human-based review process, and while the state of the art may not indicate an imminent replacement, developments are fast-moving and future experiments may firmly come out in favor of AI reviewing.

Moving beyond reviewing, Sarker et al. offer a general framework for thinking about the possibilities of human-AI configurations by representing a 2×2 table of human versus machine dominance. The authors define dominance to include not only the allocation of tasks to either the human or the machine but also the level of control over the process and its outcomes. Thus, they describe the combination of human and machine dominance as a "complementarity of human and AI capabilities being synergistically harnessed to implement a reimagined review process" and recommend reaching it

by first experimenting with configurations of human dominance and machine subordinance.

In their discussion of developing theory through human-gAI collaboration, Jarvenpaa and Klein examine both the delegation of tasks to AI and the transfer of control as interrelated decisions. The delegation of tasks is limited by what the machine can do effectively. The authors assume, for example, that gAI tools "lack not just morality but also intuition, plausibility, and temporal relevance and awareness." The delegation of responsibilities and control must provide the received ingredients of developing good theory, e.g., variety, novelty, and rigor, and also consider ethical aspects such as normativity, bias, transparency, and dependency. Jarvenpaa and Klein conclude that because of the limitations of gAI in achieving these goals, delegating too much control to machines will impair theory development. By this logic, advances in gAI, say, in enabling awareness, would require a reexamination of delegation decisions.

Other researchers in related fields have distilled the issue of control to distinguish the decision of control when determining delegation. For instance, in human-centered AI research, it is assumed that the human will maintain overall control of whatever tasks are delegated to ensure that the machine always operates for the benefit of the human (Shneiderman, 2020). Even when control for delegated tasks is transferred to the machine, there is a "red button" that ensures that humans can regain control at their discretion. Maintaining overall control while delegating tasks and control of some of the subtasks (subsystems) requires certain conditions, such as understandability, as discussed in the following sections.

In any event, the delegation of responsibilities and control is a design dilemma that each journal will have to resolve for the parts and whole of its journal ecosystem. With the expected changes in the journal ecosystem—in particular, the continued progress of AI—delegation will have to be revisited periodically.

3.2 Understandability and Transparency

Any human-AI configuration requires mutual understanding (Suchman, 2007). In the context of journals, all human-machine messages must be understandable to both parties. Furthermore, understandability must be designed to satisfy multiple journal goals, including effective coordination and cooperation between humans and AI in their assumption of different roles, effective learning of reviewers and editors, and the trustworthiness of editors in automated or augmented reviews and of editorial decision makers. The transparency of processes and the understandability of recommendations and decisions are relevant to the creation, curation, and consumption processes. The SI papers primarily addressed the interpretability of AI

messages to humans in terms of ensuring effective coordination and trustworthy AI recommendations and decisions. Not all the authors agree on the necessity of transparency for all processes. For instance, Weber argues that the quality of an automated review may best be judged according to its output rather than the interpretability of its process and the recommendation's rationale. And those who desire transparency do not always know how to achieve it or if it is even possible. In any event, every journal will have to address these issues to decide, for instance, when to prohibit the use of an AI model that lacks sufficient transparency.

Ngwenyama and Rowe examine the need for transparency in the various functions of creating literature reviews and warn us that current LLM fall short on providing transparency, which could prevent the delegation of certain functions, such as creating a meta-analysis, to the machine They argue that transparency is a core value in any scientific work that is expected to explain the theoretical arguments and methodological procedures; without this, the evaluation of the argumentation, logical consistency, and completeness will be lacking. Ngwenyama and Rowe believe that current LLM and other ML tools cannot deliver the transparency needed and cannot confirm their adherence to the received epistemic practices of the field. This is a good example of how transparency, or the lack thereof, affects delegation and control.

The need for transparency in creation carries over to curation and consumption and adds to the considerations of when and how the transparency and understandability of AI models should be supplied. Sarker et al., too, maintain that transparency is needed to ensure that AI models instantiate our academic values of rigor, fairness, and access. They fear that without confirmation that these values are being upheld, the community will reject or avoid reliance on AI-based reviews. Sarker et al. go one step further to propose guidelines for designing the human-AI collaboration accordingly. In particular, they note the role of the community in insisting on explainable AI models and providing appropriate feedback, which will lead to higher trustworthiness and acceptance. Such involvement by our community may build on more general work such as the European Commission's (2019) Ethical Guidelines for Trustworthy AI. Sarker et al. also note that given extreme datafication, explainability and transparency are essential for monitoring ethical threats.

Shmueli and Ray tie transparency in the editorial process to accountability. AI augmentation can enhance editorial decision-making and, at the same time maintain human judgment and accountability. However, it must address the question of algorithm transparency to ensure conformity with expected norms and values. These expectations will vary with the type of AI technology adopted. For instance, in persuasive technology, some actions are transparent, such as suggestions by recommender systems, while others are not, such as comment moderation on social networks and deceptive interface design choices (Greene et al., 2022). Indeed, in AI augmentation, interpretability and performance may produce a trade-off between alternative AI technologies, contingent on the goal of the augmented function. For example, machine-learning text classification models that are less explainable compromise learning even though they may be more accurate, thus improving classifications but reducing trustworthiness (Te'eni et al., 2023).

The transparency and understandability dilemma is contingent on the process, goal, and AI technology, and it is complicated, involving several considerations, as demonstrated in the three papers discussed above. Each journal will tackle this dilemma according to its values, priorities, and the processes it wishes to augment or automate.

3.3   Accountability and Responsible AI

With delegation comes accountability, which is closely related to transparency. Accountability is one of the design principles of responsible AI (Gregor describes these principles in detail) and is therefore a critical element of designing AI-augmented journals. The accountability of human agents and the accountability of artificial agents are distinct yet share the basic idea that accountability is a relation between an agent and a principal such that the agent must justify its conduct to the principal, and the principal supervises, asks questions, and evaluates the agent on the basis of such justification (Novelli et al., 2023).

All the papers in this SI assume that the human EIC has ultimate control over the AI-augmented journal and, therefore, that human-to-human accountability remains as is at "the top" but changes, with varying degrees of control envisioned throughout the creation and curation processes. Journals and EICs are held accountable to the community and to other stakeholders such as the publisher; senior and associate editors are held accountable to the EIC, reviewers, and authors; and reviewers are held accountable to the editors and authors. Delegation does not remove accountability. Hence, all these human agents are accountable for the outcomes of the responsibilities they delegate. Jarvenpaa and Klein, for example, state clearly that while gAI tools can support knowledge creation by theorist teams, the theorists remain accountable for the validity of the resulting theory.

AI accountability is necessary to enable the accountability of the human principal who delegates responsibilities to the AI systems. This relates to the expectation that the designers and deployers will comply with standards and legislation to ensure the proper functioning of AI systems during their lifecycle (Novelli et al., 2023). This requires a sociotechnical design of AI systems to provide compliance, reporting,

oversight and enforcement, which enable the principal to be accountable for the validity of the products but also for the appropriateness of the process of the AI system. The implication is that AI accountability is designed not only to justify the end product of the system, e.g., a journal publication, but to justify each task delegated to the artificial agents, e.g., supporting the resolution of theoretical claims or reviewing the submission.

Shmueli and Ray present both AI-augmented and AI-driven scenarios. In the former, AI models provide predictions and recommendations to a human associate editor, who remains accountable to the EIC, reviewers and authors. The journal must make decisions based on the level of algorithm transparency, the appropriateness of the AI model, and other ethical considerations that must be upheld such as data privacy and security. Shmueli and Ray clearly state: "The human remains firmly in control and accountable: The editorial process cannot proceed without timely actions of editors and reviewers. Accountability and safety are supported by logging all AI and human actions, making them available for audit and system improvement. We contend that AI-augmented IT systems can empower human decision makers—editors, reviewers, and authors—who will still be accountable at every step."

This vision cannot materialize without the careful design of AI accountability, contingent on the particular context and values of the journal (Novelli et al., 2023). AI accountability is also contingent on the human's capability to evaluate AI outcomes and processes in order to stay in control, which may become increasingly difficult with advanced gAI, as demonstrated by Drori and Te'eni. Yet in the latter AI-driven scenario, far more autonomy is granted to multiple AI agents with potentially precarious results, and while Shmueli and Ray argue against this eventuality, they rightly recognize the potential of predatory publications adopting AI-driven practices that could be detrimental to the honest, fair, and responsible curation of published research. Accountability should be seen as part of responsible AI. Gregor's paper is devoted to the topic of responsible AI in journal publishing. Besides accountability, transparency, and human control, which we have discussed above, responsible AI should also address professional responsibility, safety and security, fairness and nondiscrimination, privacy, and the promotion of human values and legal rights. Accountability should refer to all these elements. Our discussions of the three challenges above refer to the journal (organization) level. Gregor shows how these aspects should also be examined at the levels of human-AI team, industry, and government regulation. This point underscores the need for the careful and contingent design of AI accountability. It also begs the question of scope and levels of analysis when addressing the impact of AI on the journal ecosystem, as seen in the next subsection.

3.4 Impact: The Scope, Extent, and Intensity of Change Affected by AI

In this fourth challenge, we look at AI's impact both on the journal creation, curation, and consumption processes and on the context of the journal. Future frameworks will likely refine this broad category that was fitted to integrate the ideas and the stakeholders mentioned in the SI. The journal ecosystem serves several stakeholders, including authors, reviewers, editors, professional communities, institutions, associations, and publishers. Under context, we included the infrastructure on which the journal system operates. It consists of (1) the technological infrastructure, including AI-based tools, distributed ledgers, and searching and structuring tools; (2) the epistemic infrastructure, consisting of the systems, institutions, and mechanisms supporting knowledge creation, curation and consumption; (3) the economic and legal infrastructure, which includes the market matching the consumers and producers of journals operating under laws and regulation; (4) the social infrastructure, representing the broad society impacted by the changes introduced by research; and (5) the professional community of academic authors and readers, and possibly others such as IS practitioners. Most of these subcategories are discussed in at least one of the 11 papers, although, as can be expected, the categories are not examined comprehensively. For example, none of the authors examined the implications of legal aspects.

Defining and predicting the impact of new applications of AI in the journal ecosystem is challenging for several reasons. First, it is difficult to predict the outcomes of generative technologies and platforms, by definition. Only with experience does one typically discover unforeseen effects on other technologies, people, organizations, and societies. Second, impact is framed differently by different stakeholders—editors, reviewers, and authors in the case of reviewing, and the community of readers, associations such as the AIS, publishers, and the constituents in academic promotion and tenure processes in the case of the entire journal ecosystem. Third, the relationships between AI technologies and the elements of the ecosystem are neither linear nor unidirectional; as the technologies modify processes, the modified processes trigger new technologies or adaptations to extant technologies. For all these reasons, predicting, defining, and understanding the impacts will be a continual challenge for the designers and adopters of AI in the journal ecosystem. Several opinions, nevertheless, project changes in the processes and the infrastructure.

Weber's and Kankanhalli's papers begin with the impact of AI on the review process and proceed to talk about the broader impacts such as changes in the product under review but also changes in other steps of the editorial process. Changes in the review process

may result in an unbundling of the product into elements of an issue or even elements of a paper, and changes in the editorial process may result from new patterns of delegation. For instance, the availability of automated reviews will encourage authors to review their papers before submission and correct them according to the feedback, which may result in a more efficient process and higher-quality papers. However, Drori and Te'eni caution us to examine the risks that come with the opportunities in AI-augmented reviewing. The risks, including bias, inappropriate practices, and misuse, are often hidden or discounted.

Shumeli and Ray present an analysis of both positive and negative impacts for each stakeholder they identify in the ecosystem. Interestingly, they advocate a future in which AI augments human capabilities in their different roles and supports academic (human) discourse.

Several papers advocate or predict major changes in infrastructures. Watson et al. argue for the digitization of published causal models to create a graph database that enables causal knowledge analytics. They maintain that LLMs cannot extract knowledge with sufficient accuracy to enable the computation of causal knowledge metrics. Furthermore, digitization, a form of tagging, facilitates combining AI and knowledge graphs to digitally transform scholarship in order to overcome current scope and scale boundaries on the curation and creation of knowledge. Ngwenyama and Rowe claim that the epistemic infrastructure must be understood in its organizational and social context. They warn of the sometimes hidden risks resulting from AI's inability to fit social and organizational impositions. Moreover, organizational pressures and technological temptations may overshadow the risks of failing to meet social and organizational requirements, such as the expected transparency and understandability of the research method, a phenomenon the authors call false consciousness (Ngwenyama et al., 2023). This dangerous misjudgment of risk versus opportunity requires that special attention be paid to revealing the negative impacts of AI. Finally, Yoo directly links AI-based changes in the epistemic infrastructure to a revolution in the creation, curation, and consumption of knowledge. He identifies two directions that will potentially bring about such dramatic changes—namely, decentralization and temporary binding—and demonstrates the role that generative AI may play and its implications on the form and function of journals. For instance, in the traditional role of journals as both gatekeepers and developers of knowledge, Yoo sees a strong shift toward knowledge development.

Focusing on the economic and social infrastructure of the academic publishing ecosystem, Avital calls for the decentralization and democratization of journals relying on blockchain and AI technologies to transform the socioeconomic infrastructure into a token-based market for managing the relationships and equity distribution among editors, reviewers, researchers, and readers. He examines how the scientific community can utilize emerging technologies to transform the institutional logic of journals and harness AI and token-driven organizational governance to nurture a participative scholarship culture and overcome a market failure that stifles knowledge dissemination and leaves commercial publishers with unsustainable leverage.

These transformations in economic, social, and epistemic infrastructures, enable and intensify potential digital transformations in the creation, curation, and consumption processes described above, but at the same time, they are also affected by the processes. The constant mutual effects of process and infrastructure create a dynamic journal ecosystem in which journals will have to frequently reinvent themselves to cope with the rapid changes in technology and ways that we produce and consume knowledge.

Clearly, journals will have to painstakingly debate the impact of AI on the journal ecosystem and their specific journal with its particular epistemic and social context. It will require analyses at multiple levels and the rippling effects from one level of analysis to another, as demonstrated in this SI. While the papers demonstrate directions and pitfalls in assessing impacts, there is, unfortunately, much to explore in how to measure impacts, especially negative impacts.

In sum, we discuss four thematic challenges that should be considered by journals when contemplating the introduction of AI into their management and operations. Taken together, these challenges demonstrate the complexity and importance of considering the capabilities, opportunities, and risks involved when employing AI technologies. The diverse opinions about these challenges presented in the SI papers point toward pressing research issues that are expected to become more important as we learn more about the impacts of various AI technologies and as new technologies develop. We chose these four challenges, as they effectively encompass most of the opinions in this SI, with the understanding that this is an incomplete list of challenges we face. We hope this presentation of four challenges motivates readers to add other aspects critical to decision-making and future research.

4 Conclusion

Now what? A call to action, or a call to contemplation, conversation, design, and assessment?

We see a pressing need for both. Action, in the form of the indiscriminate introduction of AI to the creation, curation, and consumption of knowledge, can have irreversible perilous results on the journal ecosystem—

and dare we say to the epistemic basis of knowledge ecosystems. Yet lack of action is not really an option. There is already a rush that will undoubtedly accelerate to introduce different forms of AI across unscrupulous disreputable journals in many disciplines without regard to the opinions expressed in this SI. And like environmental ecosystems in which one water supply is tainted, the effects of a tainted knowledge supply can damage us all.

The actions that we need right now will be defensive— the creation of filters, detectors, blockers, and other forms of technology and procedures to impede the advance of predatory AI. This was needed when predatory journals and paper mills began to clog our libraries and will be needed tenfold now that much of that activity can likely become AI driven and too easy. By studying these actions, as information systems researchers and professionals, we will undoubtedly learn valuable lessons regarding what should and should not be done. What forms of delegation and control are most effective? What types of understandability and transparency can be harnessed? How does the need for accountability and responsibility change our organizational structures and processes? And what forms of external impact create lasting value rather than momentary whimsy followed by chaos?

In parallel to those actions must come the contemplation, conversation, design, and assessment that the authors in this SI have collectively called for. Alongside this must be the responsible introduction of AI that will allow our human journal teams to transparently introduce new efficiencies to our journal ecosystems without compromising fairness, quality, and impact. With a heightened awareness of the principles and processes, checks and balances, risks and benefits raised in this SI, we must do what IS leaders do best—raise the lighthouse to guide our peers and other disciplines, showing the way to safely embrace new, poorly understood, risky, highly impactful technology. When Phillip Ein-Dor took the

first visionary steps to create JAIS as an online journal of high quality and repute, there were already multiple online predatory journals exploiting the new technology, which we are still battling today almost 25 years on. But that didn't stop Phillip from taking those first steps—and it shouldn't stop us.

Acknowledgments

References

Ågerfalk, P. J. (2020). Artificial intelligence as digital agency. *European Journal of Information Systems*, *29*(1), 1-8.

AIS. (2014). *Code of Research Conduct.* https://aisnet.org/page/AdmBullCResearchCond

Checco, A., Bracciale, L., Loreti, P., Pinfield, S., & Bianchi, G. (2021). AI-assisted peer review, *Humanities and Social Sciences Communications, 8*(1), Article 25.

COPE. (2023). *Authorship and AI tools: COPE position statement.* https://publicationethics.org/cope-position-statements/ai-author

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kumar Kar, A., Baabdulla, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Ahmad Albashrawi, M., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., Carter, L., … Wright, R. (2023). "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy, *International Journal of Information Management* (71), Article 102642.

European Commission (2019). *Ethics guidelines for trustworthy AI.* https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

Eva, A. M., Bollen, J., Zuidema, W., & Bockting, C. L. (2023). ChatGPT: Five priorities for research. *Nature, 614*, 224-226.

Greene, T., Martens, D., & Shmueli, G. (2022). Barriers to academic data science research in the new realm of algorithmic behaviour modification by digital platforms. *Nature Machine Intelligence*, *4*(4), 323-330.

Kumar, A., Ghosal, T., Bhattacharjee, S., and Ekbal, A. (2023). Towards automated meta- review generation via an NLP/ML pipeline in different stages of the scholarly peer review process, *International Journal on Digital Libraries*. https://doi.org/10.1007/s00799-023-00359-0.

Ngwenyama O., Rowe F., Klein S., & Henrikssen. H. (2023). The open prison of the big data revolution: False consciousness, faustian bargains and digital entrapment, *Information Systems Research*. Advance online publication. https://www.doi.org/10.1287/isre.2020.0588

Nature. (2023). *Tools such as ChatGPT threaten transparent science; here are our ground rules for their use* [editorial]. https://www.nature.com/articles/d41586-023-00191-1

Novelli, C., Taddeo, M., & Floridi, L. (2023). Accountability in artificial intelligence: what it is and how it works. AI *& SOCIETY*. https://doi.org/10.1007/s00146-023-01635-y

Pradhan, T., Bhatia, C., Kumar, P., and Pal, S. (2021). A deep neural architecture based meta-review generation and final decision prediction of a scholarly article, *Neurocomputing*, *428*, pp. 218-238.

Shmueli, G., Maria Colosimo, B., Martens, D., Padman, R., Saar-Tsechansky, M., R. Liu Sheng, O., Street, W. N., & Tsui, K.-L. (2023). How can IJDS authors, reviewers, and editors use (and misuse) generative AI? *INFORMS Journal on Data Science*, *2*(1), 1-9.

Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, *36*(6), 495-504.

Steingard, D., & Linacre, S. (2023). Transforming academic journal assessment from "quality" to "impact": A case study of the SDG impact intensity academic journal rating artificial intelligence system. In C. Hauser & W. Amann (Eds.), *The future of responsible management education: University leadership and the digital transformation challenge*. (pp. 317-356). Springer.

Stokel-Walker, C. (2023). ChatGPT listed as author on research papers: Many scientists disapprove *Nature, 613*, 620-621.

Suchman, L. A. (2007). *Human-machine reconfigurations: Plans and situated actions*. Cambridge University Press.

Susarla, A., Gopal, R., Thatcher, J. B., & Sarker, S. (2023). The Janus effect of generative AI: Charting the path for responsible conduct of scholarly activities in information systems. *Information Systems Research, 34*(2), 399-408.

Te'eni, D., Zagalsky, A., Yahav, I., Schwartz, D.G., Silverman, G., Cohen, D., Mann,Y. & Lewinsky, D. (2023). Reciprocal human-machine learning: A theory and an instantiation for the case of message classification. *Management Science*. Advance online publication. https://doi.org/10.1287/mnsc.2022.03518

Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science, 379*(6630), 313.

Yuan, W., Liu, P., and Neubig, G. (2022). Can we automate scientific reviewing? *Journal of Artificial Intelligence Research, 75*, 171-212.

Honoring the Memory of Phillip Ein-Dor

Professor emeritus Phillip Ein-Dor of Tel Aviv University died on November 11, 2022 at the age of 88. A prominent figure in substantiating the field of information systems, Ein-Dor established our flagship *Journal of the Association for Information Systems* (JAIS) as its founding editor, led the Association for Information Systems (AIS) as its eighth president, and chartered the Israel chapter of AIS.

First and foremost, Phillip Ein-Dor will be remembered for his kindness, friendliness, wisdom, and modesty, and for his quiet talk and bold impact on individual members of the IS community and on institutional initiatives. He never ceased to help students and colleagues with advice, guidance, and pragmatic support in their career development, and never ceased to be deeply involved in the profession, in local and global academic affairs, and in national education and governance of the profession. Phillip was a dear friend and mentor to so many of us.

During his term as AIS president (2002-2003), Phillip advanced the issues of globalization, transparency, and defining the nature of the IS field. Globalization initiatives included funding for surveys of the IS field in eastern Africa and the Asia-Pacific region and scheduling the first AMCIS meetings outside the US and Canada starting with AMCIS 2006 in Mexico. Transparency included an improved budgeting process—in particular, the creation of a combined budget for AIS and ICIS following the affiliation of the two organizations. During his tenure, to help define the nature of the information systems field, a meeting of senior people in the field was held at ICIS in Barcelona and the Blue-Ribbon Committee was formed to examine the role of the IS field.

Phillip Ein-Dor's research interests included management information systems, information system theory, artificial intelligence—especially natural language processing and commonsense knowledge representation, the economics of computers and information systems, technology infrastructure and diffusion, internet applications, and the digital divide. He published four books, seven book chapters, and more than 50 papers on various aspects of information systems and their management. His most frequently cited works are papers concerning the relationship between organizational context and management information systems and a book proposing a paradigm for management information systems, all co-authored with Eli Segev.

Upon retirement in 2003, he served as professor of information systems at Tel Aviv University where he remained professor emeritus. Continuing to influence the field even then, Phillip was a faculty member of the Academic College of Tel Aviv-Yaffo, where he established an undergraduate program in information systems and worked on establishing an MBA program. Phillip directed numerous master's and PhD theses supporting generations of scholars. He held visiting appointments at leading departments on four continents, including New York University, Claremont Graduate University, the Naval Postgraduate School, The University of Capetown, the National University of Singapore, City University of Hong Kong, and ESSEC Business School.

Phillip was chosen as an AIS Fellow in 2000 and received the LEO Award in 2006. In Israel, he was recognized as one of the founders of the IS discipline in the country and in 2010 received an award from the Israel Chamber of Information System Analysts for "Primacy and Contribution in Founding, Establishing, and Strengthening the Information Systems Area in Academe." Additional awards include the incumbency of a Hui Sun Chair at National Chung Hsing University, Taichung, Taiwan (2002) and an award from the Fourth Mediterranean Conference on Information Systems in "Recognition of Outstanding Contribution to the Field of Information Systems in the Mediterranean Region." He will always be remembered for his incredible impact encompassing diverse research contributions in human-computer interaction, AI, and the philosophy, management, and economics of IS.

About the Authors

David G. Schwartz is a professor of information systems and former vice chairman at the Graduate School of Business of Bar-Ilan University, Israel, where he chairs the Doctoral Program. David currently studies digital health interventions and human-AI configurations. He has published over 170 articles and editorials which have appeared in publications such as *Information Systems Research, IEEE Intelligent Systems, IEEE Pervasive Computing, ACM Computing Surveys, Decision Support Systems, Management Science, Internet Research,* and *JMIR mHealth & uHealth*. David has held multiple visiting scholar positions: Columbia University, USA (2004); Monash University, Australia (2007-2008); Victoria University, Australia (2016); National Tsing Hua University, Taiwan (2017; 2020-2021); Erskine Fellow at the University of Canterbury, New Zealand (2016); and visiting professor at Renmin University of China, Beijing (2018). David served as editor-in-chief of the journal *Internet Research* for over a decade and has served as an associate editor for *European Journal of Information Systems* and *JMIR mHealth & uHealth*.

Dov Te'eni is a professor emeritus at Tel Aviv University, where he was research associate dean and held the IS Mexico Chair. He is now visiting at the Department of Middle East Studies at Bar Ilan University. Dov currently studies human-AI configurations, models of smart human-computer interaction, and knowledge sharing. He co-authored *Human-Computer Interaction for Developing Effective Organizational Systems* with Ping Zhang and Jane Carey and co-edited the *Encyclopedia of Knowledge Management* with David Schwartz, as well as other books on information systems and innovation. He has written close to two hundred papers with over seventy co-authors in journals such as *Management Science, MIS Quarterly, Organization Science, Journal of the Association for Information Systems,* and *International Journal of Human-Computer Studies*. Dov is the past president of AIS—the international Association of Information Systems, and past editor-in-chief of the *European Journal of IS*. Dov has been recognized with the AIS Fellow award (2008) and the lifetime-achievement LEO award (2015).