Bus Inf Syst Eng 67(6):771–775 (2025)
https://doi.org/10.1007/s12599-025-00972-0

EDITORIAL

Process Mindlessness: When we Lose Sight of What AI is Supposed to Improve

Wil M. P. van der Aalst · Hajo A. Reijers · Oliver Hinz · Jella Pfeiffer

Published online: 21 October 2025

1 The Productivity Trap: When AI Solves Problems it Created

The uptake of AI is amazing. Notably, recent developments in generative AI have enabled the effortless processing of text and images, the automation of simple tasks using agents, the automatic summarization of meetings, and the creation of new content. AI can accelerate operational and development processes, making them faster and more cost-effective. Most major tech companies have announced AI-related layoffs, highlighting the potential for AI to automate tasks previously performed by humans (Goldman Sachs 2025). This is also unexpectedly impacting highly skilled jobs in sectors such as software development.

W. M. P. van der Aalst (✉)
RWTH Aachen, Lehrstuhl Für Informatik 9, Ahornstr. 55, 52056 Aachen, Germany
e-mail: wvdaalst@pads.rwth-aachen.de

H. A. Reijers
Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands
e-mail: h.a.reijers@uu.nl

O. Hinz
Faculty of Economics and Business Administration, Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany
e-mail: ohinz@wiwi.uni-frankfurt.de

J. Pfeiffer
Karlsruher Institut Für Technologie (KIT), Institut Für Wirtschaftsinformatik (WIN), Kaiserstraße 89-93, 76133 Kasrlsruhe, Germany
e-mail: pfeiffer@kit.edu

AI offers enormous potential, but it also presents a range of significant and intricate problems. These problems span technical, ethical, social, and geopolitical dimensions, posing new challenges for information systems engineering. AI systems can inherit and amplify biases present in training data, leading to unfair outcomes. AI can generate highly convincing fake content (deepfakes, fake news, etc.). AI relies on massive datasets, often collected without full consent, and enables authoritarian surveillance. Last, but not least, advanced AI capabilities are primarily controlled by a few big tech companies and wealthy countries, leading to new forms of inequality.

The problems just mentioned are real and highly relevant, but the underlying assumption is that AI indeed works as intended and is able to speed up processes and save costs. In this editorial, we aim to challenge this assumption and highlight potential situations where AI may not achieve the desired efficiency and quality gains. These are situations where customers, researchers, citizens, and employees do not experience positive effects. For example, AI may generate additional work or increase stress while trying to achieve the opposite (Mirispelakotuwa et al. 2025). We refer to the application of AI without considering the overall process as "process mindlessness".

In this editorial, we highlight three problematic effects of applying AI that degrade the overall process, rather than improve it. In this way, we *blame* the inappropriate use of AI for its *bloating*, *blurring*, and *blasting* effects (see Fig. 1).

Problem 1: Bloating. The quote "If I had more time, I would write a shorter letter" is often attributed to Mark Twain and Benjamin Franklin. It is based on the earlier statement "I have made this longer than usual because I haven't got the time to make it shorter" made in 1657 by the French mathematician and philosopher Blaise Pascal.

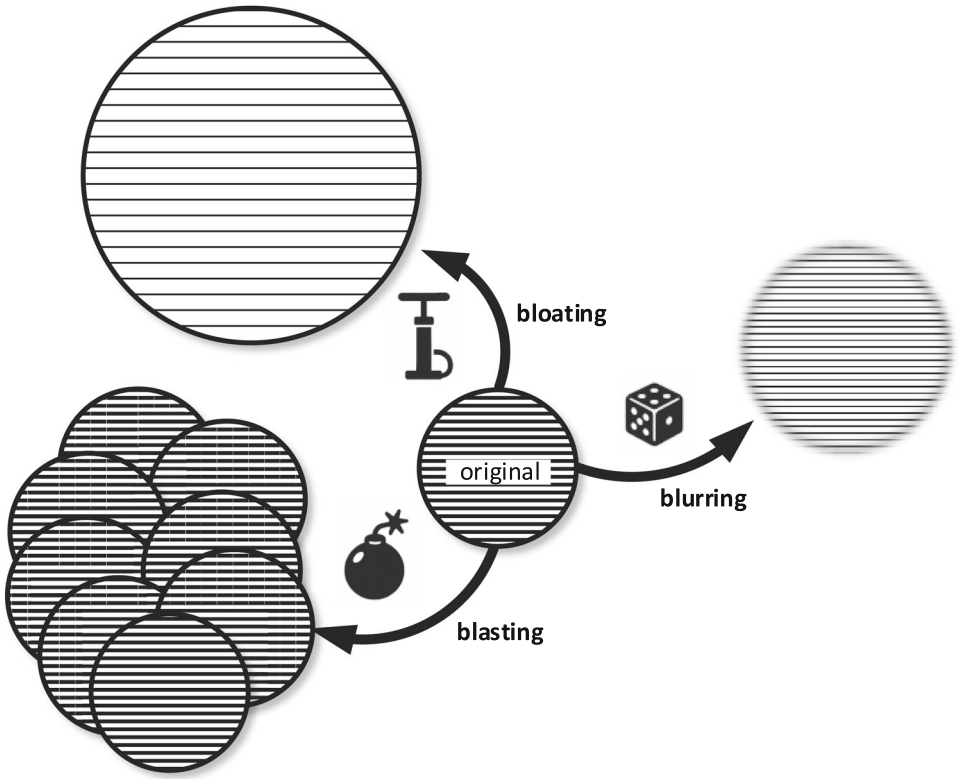

Fig. 1 Blaming AI use: the problems of bloating, blurring, and blasting

The quote points out that it takes effort to be succinct and precise. Using generative AI, it has become easier to generate large amounts of seemingly high-quality text from a prompt. Because the effort required to produce text has decreased, humans are confronted with ever-increasing amounts of text that lack diversity (Doshi and Hauser 2024). This results in "text overload" situations and stress. Any meeting can be automatically converted into a report, and based on a few hints, an extensive report can be generated using background information in the LLM. The receiving side can attempt to utilize generative AI to summarize the lengthy text produced by the sender. Both sides can try to optimize their tasks, but it seems counterproductive to first inflate text and then reduce it again. This corresponds to applying a function f on one side and applying the inverse function f^{-1} on the other side to obtain the original intent, i.e., $f^{-1}(f(x)) \approx x$. At best, bloating can be neutralized by the counter-use of AI; more often than not, the bloating effect of AI overloads process participants and wastes resources.

Problem 2: Blurring. Organizations have large amounts of structured information. Consider, for example, an information system that stores data about orders, customers, suppliers, products, invoices, deliveries, and payments, among other things. This is structured information having well-defined semantics. This is complemented by unstructured information in e-mails or meeting notes. The goal is to make informed decisions and take effective actions based on all available information. These decisions and actions tend to be structured, for example, changing suppliers or rejecting a request. Due to the remarkable capabilities of generative AI, it is tempting to transform structured information into text and then use this to generate new text (e.g., a recommendation) to inform decisions or take actions. The problem is that structured information is translated into text, thereby fuzzifying the input. The output is not based on a fixed rule or computation. Instead, the response is stochastic in nature. The blurring effect of AI obfuscates information and deteriorates decision-making.

Problem 3: Blasting. AI-based systems scale effortlessly and thus "blast away" traditional boundaries. While training or fine-tuning a model can be expensive, applying the model is typically fast and cost-free. Digital communication amplifies this effect. Unlike physical communication, it comes with virtually no cost, no time delay, and no capacity constraints. When sending paper documents, practical limits apply, e.g., a report must fit in an envelope, postage costs money, and delivery takes time. Similarly, human involvement imposes natural limits. For instance, a municipal worker might be able to send 20 personalized letters per day, thereby capping the number of interactions with citizens. In contrast, an AI-powered bot can produce thousands of letters per day, instantly, shifting the burden

to those on the receiving end. When humans or physical resources are part of the process, there are built-in constraints: limited capacity, tangible costs, and required time. But AI and digital tools can "blast away" these boundaries. Suddenly, there are no bottlenecks, costs approach zero, and actions occur in real-time. Yet this scalability is a double-edged sword. Citizens, customers, or employees on the receiving end may become overwhelmed – flooded with messages, tasks, or demands that were once held in check by physical or human limitations. In this light, traditional constraints, such as the size of an envelope, are not flaws but features. They help keep processes manageable. By blasting away these boundaries, AI risks creating overload and chaos. Sometimes, the very constraints that AI seeks to remove are essential to maintaining balance and focus.

It is essential to note that these problems pertain to situations where AI is functioning as intended, yet contributing to the degradation of overall processes. The threshold for using AI is low, and tasks can be accomplished with limited effort. Therefore, it is essential that we adopt a process perspective, which reminds us of the overall goals we aim to achieve and the risk of sub-optimization that can result from focusing on a single task.

2 Implications for Academia, Business, and Society

To illustrate the effects of process mindlessness across various domains, we will use a few examples.

Imagine a committee meeting where important decisions need to be made (e.g., the board of governance of a university or the selection committee evaluating research proposals). In the past, one would get an agenda and paper documents to prepare for such a meeting. People manually created these documents without using generative AI. Also, the documents had to be sent to all committee members. To allow people to read the documents, they were sent at least one week before the meeting. Ideally, the documents would fit into an envelope. The EN 13724 standard specifies the requirements to be met by postal mailboxes, and these only allow for envelopes less than 3 cm thick. 3 cm is roughly equivalent to 500 pages, assuming each sheet of paper is 0.06 mm thick. Therefore, there were clear physical limitations to the transfer of information. Those preparing for the meeting would carefully consider what to send, because actually printing the documents provides haptic feedback when things get out of control. Also, all documents were sent at once. Today, the documents for such a meeting easily exceed 1000 pages. Also, one week before, some documents are still missing and need to be distributed later. There may also be updates to documents that have already been distributed. As a result, committee members often experience stress and are unlikely to read the documents. The quality of decision-making is impacted by this. This is caused by a lack of physical constraints and bloated human-readable artifacts. Of course, committee members may use AI to handle such challenges, but clearly, the original goal of having a productive meeting and making informed, high-quality decisions has been lost. Things get "lost in translation" and stakeholders feel disconnected and alienated.

Generative AI is increasingly used by authors to assist in writing papers, such as summarizing prior work, refining language, and even generating results sections or related work based on a few keywords (Fecher et al. 2025). Authors are exploring the limits of what is allowed and possible. The current consensus seems to be that it is acceptable to use AI to improve presentation, but not to generate content. However, this is impossible to verify as generative AI continues to improve. Although reviewers are not supposed to use AI for reviewing (for example, to summarize or draft parts of their reviews), this is also difficult to verify. Each of these local applications of AI may seem helpful or efficient. However, the end-to-end process suffers. Authors may submit papers with inflated or repetitive content because it costs almost nothing to generate 16 pages of grammatically correct text. Reviewers now face more submissions, greater verbosity, and the temptation to skim or offload reading to an LLM. This creates a loop in which AI-generated content is reviewed by AI-based reviewers, leaving humans disconnected (both in terms of depth and responsibility). The original goal of ensuring scientific quality and advancing knowledge becomes lost in the process. To advance science, it is vital that papers are actually read and discussed.

Education is suffering from the same problem (Dong et al. 2025). Many students now rely heavily on generative AI to complete assignments, generate essays, solve math problems, or even simulate programming projects. The entry barrier is low: a few prompts can yield a grammatically correct, seemingly well-structured result in minutes. As a result, students may submit work that was mostly or entirely generated by AI without a proper understanding of the material. Consequently, these students learn little and become increasingly dependent on AI. Although this is unavoidable in many situations, teachers and professors are expected to uphold traditional roles, i.e., assessing the quality of student work, providing feedback, detecting plagiarism, and assigning grades. The use of generative AI by educators is often prohibited, particularly for grading or providing individualized feedback. This creates an asymmetry: students use AI to produce content, while teachers are not allowed to use AI to evaluate it. Considering the goals and affordability of education, the reverse may be desirable.

These examples show that solving local problems (e.g., "make writing easier" or "make reviewing faster") without considering the entire process and the original goals can degrade the system as a whole. This can also be seen in larger organizations.

3 Possible Solutions

This editorial aims to reflect on the problem of "process mindlessness" generated by new forms of AI. The rise of generative AI and the resulting information overload can be understood as a rise in entropy within organizational and societal systems. There are three obvious recommendations we can offer to turn that tide.

First of all, there should be *costs* associated with *communication*. In other words, if it is not worth being printed, do not send it. Due to digitalization and generative AI, it is too easy to produce and distribute large amounts of human-readable information. Information overload is not a new problem (Bawden and Robinson 2020). Alvin Toffler coined the term as early as 1970 (Toffler 1970). Additionally, several proposals have been made to incorporate costs into communication (e.g., electronic stamps, proof-of-work schemes, and micropayments). These attempts were not successful (just look at your own mailbox). However, it is good to creatively think about such solutions given the recent developments in AI. In organizations and communities, one could introduce quotas; rewards are due for data that are succinct, accurate, and actionable.

Second, we should rely on *structured* information and automate processes *properly* (instead of translating structured information to text and vice versa). If a chatbot can do it, it can also be adequately automated. Using generative AI may increase entropy in systems and processes. Entropy measures the uncertainty or disorder in a system. In this context, it is not just about the volume of content being produced, but also the dilution of meaning, loss of structure, and breakdown of the signal-to-noise ratio that occur when structured information is converted into text. Translating structured data into unstructured text increases entropy by destroying known relationships and injecting probabilistic noise.

Lastly, the use of AI, particularly generative AI, should be complemented with *first-principles thinking*, considering the end-to-end process and the original goals. Questions to be asked before using AI to automate a task are: How are other tasks being influenced, particularly by the output that AI creates? How are the people working in the processes affected? Does the local "upgrade" improve the process as a whole? In this sense, process thinking is needed to constrain the mindless use of AI, cf. (Grisold et al. 2024).

4 Developments of BISE Over the Last Decade

This is also the last editorial of the first author, Wil van der Aalst, in his role as a Vice-Editor-in-Chief (VEiC) of *Business & Information Systems Engineering* (BISE). It has been a great honor and privilege to serve in this capacity since 2015, coinciding with a pivotal moment in the journal's history: the transition from *Wirtschaftsinformatik*, the flagship journal for the German-speaking business informatics community, to the more international BISE. The last bilingual issue appeared in December 2014, and to support this transformation, the editorial board was made more international. I was invited to join as VEiC not only due to my Dutch background but also to help strengthen the journal's ties to the global Business Process Management (BPM) community.

BISE was founded in 1959 and now boasts a 66-year legacy, reflecting the remarkable evolution of information technology (van der Aalst et al. 2018). During my tenure, I had the pleasure of working with a series of outstanding colleagues: first with Martin Bichler as Editor-in-Chief (EiC) and Armin Heinzl as fellow VEiC; then with Christof Weinhardt as EiC and Oliver Hinz as VEiC; and most recently with Oliver Hinz as EiC and Jella Pfeiffer as VEiC. These collaborations have been immensely rewarding, and I could have easily continued in this role. However, I strongly believe in the importance of renewal within editorial teams, as it creates space for emerging talent and new ideas. I am therefore delighted that my successor is Hajo Reijers, a long-time friend and respected BPM scholar. I am confident that the journal is in excellent hands.

Looking back, the journal's progress over the past decade has been extraordinary. The two-year impact factor has increased to 10.4, with a five-year impact factor of 10.8 (as of 2024). Article downloads have increased substantially, underscoring BISE's growing visibility and relevance. The journal now ranks ahead of most other information systems (IS) journals in terms of impact. Nonetheless, there is still room to improve geographic diversity, an understandable challenge given Europe's leadership in areas like BPM, but one we must proactively address by engaging broader communities both topically and geographically.

Thematically, the journal has also evolved. While modeling and modeling languages once dominated, recent years have seen a decisive shift toward data-driven approaches, leveraging machine learning and AI. This editorial also reflects that shift. BISE is well-positioned for the coming decade, as organizations and society navigate the opportunities and challenges of AI. We are witnessing an ongoing redistribution of work, from people to machines and algorithms. The BISE community is uniquely equipped

to guide this transformation, bridging IT systems, data, organizations, and human behavior. This includes addressing new challenges, such as the risks of process mindlessness!

Funding Open Access funding enabled and organized by Projekt DEAL.

References

Bawden D, Robinson L (2020) Information overload: an introduction. Oxford Res Encycl Polit. https://doi.org/10.1093/acrefore/9780190228637.013.1360

Dong X, Wang Z, Han S (2025) Mitigating learning burnout caused by generative artificial intelligence misuse in higher education: a case study in programming language teaching. Informatics 12(2):51. https://doi.org/10.3390/informatics12020051

Doshi AR, Hauser OP (2024) Generative AI enhances individual creativity but reduces the collective diversity of novel content. Sci Adv. https://doi.org/10.1126/sciadv.adn5290

Fecher B, Hebing M, Laufer M et al (2025) Friend or foe? Exploring the implications of large language models on the science system. AI Soc 40:447–459. https://doi.org/10.1007/s00146-023-01791-1

Grisold T, Janiesch C, Röglinger M, Wynn MT (2024) "BPM is dead, long live BPM!" – an interview with Tom Davenport. Bus Inf Syst Eng 66:639–642. https://doi.org/10.1007/s12599-024-00880-9

Mirispelakotuwa I, Syed R, Wynn MT (2025) Automation to Agitation: Unveiling RPA-Induced Technostress. International Conference on Business Process Management. Springer Nature Switzerland, Cham, pp 378–394

Goldman Sachs (2025) How will AI affect the global workforce? https://www.goldmansachs.com/insights/articles/how-will-ai-affect-the-global-workforce. Accessed 12 Oct 2025

Toffler A (1970) Future shock. Random House

van der Aalst WMP, Becker J, Bichler M et al (2018) Views on the past, present, and future of business and information systems engineering. Bus Inf Syst Eng 60:443–477. https://doi.org/10.1007/s12599-018-0561-1