

Beyond Tokenmaxxing

David Mortimore^{1,2}[0000-0001-5228-4971], Mayur Patil¹[0009-0003-1154-3326], and Jenna Hillhouse³[0009-0004-7523-4909]

¹ Naval Undersea Warfare Center Division, Keyport, Keyport WA 98345, USA
{david.b.mortimore.civ, mayur.m.patil.civ}@us.navy.mil

² Carnegie Mellon University, Pittsburgh PA 15213, USA

³ Naval Postgraduate School, Monterey CA 93943, USA
jlhilllo@nps.edu

Abstract. The pursuit of greater performance and goal attainment motivates organizations to invest in and expand use of artificial intelligence (AI) agents to accomplish work. Modern AI agents often possess capabilities that are on par with and, in many cases, exceed those of human counterparts. However, AI agents often generate negligible performance gains and returns on investment (ROIs). Existing organizational governance models might be problematic for hybrid organizations comprising human and AI agents: such models inadequately consider the interpretative nature of task performance and AI agent effects on cognitive slack. Consequently, AI agent contributions to organizational performance and ROI in AI agents suffer. This paper describes a novel governance model for hybrid organizations that is informing an ongoing study to enhance mission accomplishment of a national level research sponsor and AI agent design. This model can likely strengthen performance, human agency, organizational resilience, and enhance ROIs in AI agents. This paper recommends that future work refine this model, assess its relevance more broadly, and compare the performance it generates with other governance models.

Keywords: Human-AI Teams, Agentic AI, Task Design, Organizational Performance

1 Introduction

Organizations are increasingly incorporating artificial intelligence (AI) agents into their operations in the pursuit of greater performance and goal attainment, necessitating that scholars and practitioners revisit traditional governance structures. AI agents possess capabilities that oftentimes exceed those of humans, notionally enabling them to contribute meaningfully to and strengthen organizational performance [1]–[3]. Systematically incorporating AI agents into human-centric organizations has significant governance implications [4]: theoretical and practical frameworks for designing organizations and accomplishing work generally limit AI agent contributions to group performance in hybrid organizations comprising human and AI agents [1], [5], [6].

Problematically, organizations are generally not reaping envisioned benefits from investments in increasingly capable AI agents. A 2025 MIT NANDA study of enterprise generative AI (GenAI) adoption highlighted this problem. This study found that approximately 95% of enterprise GenAI initiatives generated little or no measurable

impact on profit-and-loss performance based on a systematic review of more than 300 publicly disclosed AI initiatives, interviews with representatives from 52 organizations, and survey responses from 153 senior managers [7]. Similarly, KPMG's *Global AI Pulse* found that only 8% of surveyed organizations reported positive returns on investment (ROIs) from AI initiatives despite 95% having a formal AI strategy [8]. Organizations appear to be doubling down on AI agents despite such lackluster results. In a survey of 1,854 executives, [10] found that 85% of organizations had increased AI investments during the preceding year and 91% planned further increases regardless of ROI uncertainty. Further, [10] found that organizations typically anticipated generating a satisfactory ROI from such investments within two to four years, substantially longer than the usual seven-to-twelve-month payback period for similar investments.

Amazon and Meta, two of the leading global AI institutions arguably, generated similar lackluster results. Amazon discontinued an initiative intended to encourage use of a crude AI agent by its software developers because of significant cost increases, negligible productivity gains, and attempts to manipulate results through tokenmaxxing [10], [11]. Meta experienced similar ROI and productivity issues with crude AI agents, raising questions about how to design and manage hybrid organizations [12]. Amazon and Meta experienced additional governance issues when AI agents took inappropriate actions that disrupted Amazon services [13], [14] and exposed sensitive Meta company and user information to employees [15].

Changing organizational roles and displacing human employees with AI agents are additional governance issues needing attention. Approximately 41% of employers plan to shift work from human to AI agents, changing both the nature of work that humans perform and number of human employees [16]. Such changes are likely to affect the expertise possessed by human agents significantly and reduce agency [17]. Notably, critical thinking and reasoning skills will increasingly suffer, especially among knowledge workers, with humans attempting to assess AI agent performance and outputs without adequate understanding [18], [19]. Existing approaches to organizational design and management further increase threats to human agency, work quality, and employment opportunities [20], [21], affecting opportunities for humans to develop relevant expertise, exercise judgement, and maintain meaningful agency.

These and similar results effectively demand a new governance approach for hybrid organizations. This paper describes such an approach that, although in its early development stages, emphasizes two crucial design factors: the degree of interpretation generally needed to accomplish work and degree to which AI agents recover cognitive slack for human agents. This approach has two important benefits beyond notionally enhancing performance and goal attainment: it strengthens the roles of human agents and organizational resilience, which are increasingly important in more complex environments.

Organizational governance is described first as mainly a design effort in which task design translates organizational strategy to work accomplishment and addresses limitations on human performance. Then, the interpretation and cognitive slack (ICS) model and two key features that distinguish it from other models are summarized. An investment decision-making scenario that serves as the context for an ongoing study and the design of the associated AI agent are described next. General implications of

the ICS model are then discussed before recommendations for future research are offered.

2 Organizational Governance

Organizational governance entails purposefully designing and managing work accomplishment, performance, and goal attainment for stakeholder benefit; stronger governance approaches generally result in greater performance. Organizational tasks reify strategies for goal attainment and break down technological transformation processes into reasonably scoped and associated work components [22], [23]. Task design, therefore, involves intentionally constructing what work is performed by which group members and how they apply their expertise through relevant techniques and methods using hardware, software, and other tools [22], [24]. The tight coupling between task design and organizational performance means that strengthening governance of hybrid organizations depends largely on the degree to which task design gainfully considers human and AI agent capabilities and perspectives [23], [25].

Limitations on human performance exist and organizations provide members with the means to overcome such limitations in the pursuit of greater performance and shared goals. Cognitive, physiological, and temporal limitations bound what humans can do and accomplish [26]–[28]. Organizations generally overcome such limitations through task design, divisions of labor, and cooperative task performance [27], [29]. Nevertheless, limitations on human performance extend to and often bound organizational performance. The performance of AI agents is not necessarily limited in the same ways as humans: modern AI agents can perform work on par with and, in some cases, more productively than some humans [1]–[3]. Systematically incorporating AI agents into organizational design, operations, and management offers a means to generate greater overall performance and goal attainment [1], [25].

3 Interpretation and Cognitive Slack Governance Model

Two factors distinguish the ICS governance model and our approach to designing and managing hybrid organizations from others. The extent to which a task involves interpretation is first. *Interpretation* involves assigning contextual meaning to observed phenomena, including application of moral, ethical, pragmatic, organizational, and societal frameworks [31]–[33]. *Inference*, in contrast, refers to the process of drawing conclusions generally from a set of premises in a logical manner [31]–[33]. More complex organizational tasks, such as strategic planning, managing capital investments, deciding between societal and organizational priorities, and responding to significant and uncommon environmental changes typically involve a greater degree of interpretation than simpler and more inference-centric tasks.

While modern AI agents often possess inference capabilities equivalent to and, in some cases, greater than some humans, AI agents generally lack the ability to interpret phenomena in contextually grounded ways as humans do – a crucial design consideration. Hybrid organization governance models that intentionally consider and assign

more interpretation-centric work to humans with AI agents performing simpler and more inference-centric work are likely to outperform other models. Importantly, performing more complex tasks tends to enhance expertise and organizational standing, thereby strengthening the perceived roles of humans in hybrid organizations, consistent with findings regarding the types of work that humans prefer [33].

The extent to which AI agents return cognitive slack to humans is second. Decision-making robustness distinguishes organizational performance and goal attainment and depends significantly on effective employment of expertise and information [27], [34], [35]. *Cognitive slack* represents the dynamic amount of thinking time, or “white space,” available to individuals and organizations, and conditions decision-making robustness and, consequently, performance [36]. Greater cognitive slack generally enhances organizational decision-making robustness and goal attainment, especially when task completion and goal attainment involve greater degrees of interpretation [34], [35], [37]. Organizational governance models that inadequately consider the effects of AI agents on cognitive slack are likely to generate less optimal performance [7]–[12]. In other words, governance models that purposefully employ AI agents to recover cognitive slack are likely to outperform models that do not.

3.1 Investment Decision-Making Scenario

Research and development (R&D) investment programs provide organizations with opportunities to transform how they accomplish work, strengthen performance, and attain goals. Organizations are resource-limited and deciding between investment options is typically a significant, complex, and costly task with far-reaching effects [27], [38]. The robustness of R&D investment decision-making depends significantly on interpretation, including the extent to which program managers and teams develop and maintain awareness of previous and existing organizational investments and associated performance [27], [34], [35]. Enhancing R&D investment program performance while returning cognitive slack to the program manager and team provides the context for a study underway, including AI agent design.

Sense-making is essential to organizational investment program performance [37]. In the study underway, the program manager and team—all human agents—must develop and maintain shared awareness and assess the relevance of nearly 2,600 fundamental research studies funded by a national level R&D sponsor through dozens of distributed program managers leading differentiated investment programs. This awareness informs decision-making regarding existing and future robotics and autonomous systems research investments. Additionally, the program manager is expected to consider previous investments when managing and restructuring their investment portfolio. This sense-making task is costly: it consumes significant amounts of limited resources and less informed decision-making outcomes can have severe financial, operational, and strategic repercussions. The nature of this task, especially the need to gather, synthesize, and make sense of large volumes of information, means it is well-suited for exploring the ICS model, including how it informs AI agent design.

3.2 AI Agent Design

Enhancing performance of interpretation-centric work by the program manager and team and recovering cognitive slack drove AI agent design and development. The AI agent calculates the likelihood that each of the roughly 2,600 annual studies, describes potentially relevant fundamental research using inductive and deductive inference and graphically displays the aggregated and stratified results. The program manager and team can use recovered cognitive slack to assess the results and relevance of individual studies more fully.

To accomplish this work, the AI agent design includes a multi-layer architecture combining natural language processing and unsupervised machine learning algorithms. The main layer uses local, domain-specific sentence transformers [39] to convert unstructured text from investment program reports (IPRs) into embeddings for similarity over keyword-specific searches. Hierarchical Density-Based Spatial Clustering of Applications with Noise [40] is then applied to these IPR embeddings to identify relevant IPR groups while isolating noise. A labeling router then uses Term Frequency-Inverse Document Frequency [41], probabilistic Latent Dirichlet Allocation [41], and large language models based on user choice to create labels for the IPR groups. This design also provides a continual learning capability with human-in-the-loop active learning technique by storing user “liked” IPRs for future reference [42]. Together, these computational models function as an engine for semantic triage.

The AI agent ingests data from IPRs, including abstracts, research methodology descriptions, scholarly publication records, research/author information, study goals, accomplishments, training opportunities, and other textual descriptions, then projects the gathered data onto a unified semantic space. The AI agent then calculates the (cosine) similarity of all IPRs [43], using a “target” embedding based on user prompts or a centroid of user-affirmed exemplars. Next, the AI agent creates a visual and an interactive analysis using a “dartboard plot” based on an IPR’s syntactic and semantic proximity to the target embedding; see Fig. 1. Further categories and tiers, displayed as rings in Fig. 1, are created based on user-affirmed IPRs, semantic matching, and machine learning-derived recommendations.

This design reifies the ICS model. The AI agent performs inference-centric work, taking advantage of its computational power, grounded in program manager-affirmed observations (data points). This recovers cognitive slack that the program manager and team would have otherwise consumed developing and maintaining awareness of thousands of projects administered by dozens of program managers. The recovered cognitive slack can be used for interpretative work—crucial to goal attainment—comprehensively assessing investment project relevance, implications to existing and future investments, and deciding to invest or divest of projects based on dynamic environments and stakeholder priorities. The recovered cognitive slack can also strengthen organizational resilience and reliability [44], and contribute to more robust decision-making [34]–[36], enhancing the program manager’s organizational standing.




Fig. 1. The AI agent delivers an interactive visualization that a program manager can use to assess the relevance of thousands of investments. The bullseye (star) is constructed based on program manager-affirmed IPRs. Inner ring (orange) projects are the most relevant and identified using semantic and syntactic matches. Middle ring (red) and outer ring (blue) projects are identified using machine learning techniques. Projects in the outermost ring (white) likely have minimal relevance. This approach can enable the program manager and team to develop and maintain awareness of relevant investments more effectively and efficiently, and inform decision-making.

4 Discussion

Modern AI agents represent a transformative opportunity for organizational stakeholders seeking greater performance and goal attainment. Such agents are increasingly capable, and many can perform work as well as, and sometimes greater than, human counterparts [1]–[3]. Furthermore, AI agents are not limited in the same ways as human agents are [26]. Incorporating AI agents into traditionally human-centric organizations could effectively decouple organizational and human performance, thereby enabling organizations to generate greater results and stakeholder benefits than they might otherwise. Despite significant investments in AI agents, organizational performance gains oftentimes fail to meet expectations [7]–[15]. Furthermore, humans can “game the system” and assign less meaningful work to AI agents [10]–[12], further diminishing performance. So, what is the problem?

Governance models that inadequately consider the interpretative nature of tasks and effects that AI agents have on the cognitive slack of human counterparts appear

problematic. Human agents generally prefer to perform more complex and less routinized tasks that enhance their expertise – work that typically involves greater interpretation capabilities. AI agents tend to excel at inference-centric work, especially work that demands greater computational power. This distinction is crucial. Organizational governance models that inadequately consider this distinction risk assigning interpretation-centric work to AI agents for which they are generally ill-suited and disincentivizing human counterparts.

Additionally, governance models generally do not consider the effects that AI agents can have on cognitive slack, which represents a limited resource that drives decision-making robustness and, consequently, organizational performance [34]–[36]. Instead, organizational governance models tend to focus on less meaningful measures, such as the frequency of interactions with an AI agent and token usage [7]–[15]. Such measures are tenuously connected to organizational productivity and can incentivize human behaviors that increase organizational costs while diminishing productivity.

The ICS governance model puts these two crucial factors at the heart of hybrid organization design and management. This approach is also likely to generate secondary effects that further strengthen performance and goal attainment. The ICS model notionally strengthens human roles and standing in hybrid organizations and, by extension, agency and prestige, incentivizing humans to work with AI agent counterparts. Moreover, the ICS model can also strengthen organizational resilience and reliability [44], which are key to mission accomplishment in more complex environments.

Three limitations on the existing ICS model relevance are especially important. First, the supporting study is ongoing with comparative experimentation scheduled for later this year. Empirical findings are needed to more thoroughly assess theoretical arguments regarding ICS model usefulness. Second, approaches to distinguishing interpretation- and inference-centric work, as well as the degree to which a task involves interpretation capabilities are absent. Scholars and practitioners would likely find heuristic and other indicators useful in considering and using the ICS model. Third, the ongoing study involves a single investment program manager and team, limiting generalizability. Community members might consider ICS model merits more seriously if the scope of the study were broader.

Significant investment in hybrid organizations is likely to continue, making further exploration of the ICS model worthwhile. Future research should explore frameworks for differentiating interpretation- and inference-based work in the context of hybrid organizations and the relative capabilities of human and AI agents, systematically extend the scope of the ongoing study, and construct frameworks for assessing cognitive slack availability and adequacy.

5 Conclusion

Academic, commercial, non-profit, and government sector organizations continue investing significantly in AI agents to enhance performance and goal attainment. Limitations on human performance generally bound the extent to which organizations accomplish their missions and generate stakeholder benefits. Hybrid organizations appear

to provide a structure for moving beyond human-centric bounds, although common governance models have been unable to take full advantage of human and AI agent capabilities overall. This paper described the ICS governance model in the context of national level fundamental research investment decision-making and how it informed AI agent design and output. The ICS model and its emphasis on interpretation capabilities and recovery of cognitive slack can benefit scholars and practitioners alike. Moreover, the ICS model should enhance human agency and organizational resilience, which are increasingly important to modern organizations.

Acknowledgments. This study was funded by the Office of Naval Research Cooperative Autonomous Swarm Technology and In-house Laboratory Independent Research programs.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Burton, R. M., Obel, B., Håkansson, D. D.: Expanding the organization design space: The emergence of AI robot bosses. *J. Organ. Des.* (2023). <https://doi.org/10.1007/s41469-023-00155-9>
2. Haefner, N., Wincent, J., Parida, V., Gassmann, O.: Artificial intelligence and innovation management: A review, framework, and research agenda. *Technol. Forecasting & Soc. Change*, **162** (2021). <https://doi.org/10.1016/j.techfore.2020.120392>
3. Huang, M.-H., Rust, R. T.: Artificial intelligence in service. *J. Service Res.*, **21**(2), 155–172 (2018). <https://doi.org/10.1177/1094670517752459>
4. Bevir, M.: A reply to critics. *Int. J. of Organ. Theory & Behav.* **19**(3), 389–40 (2016). <https://doi.org/10.1108/IJOTB-19-03-2016-B007>
5. Puranam, P.: Human–AI collaborative decision-making as an organization design problem. *J Organ. Des.* **10**, 75–80 (2021). <https://doi.org/10.1007/s41469-021-00095-2>
6. Malone, T., Crowston, K.: Toward an interdisciplinary theory of coordination [MIT Sloan Working Paper #3294–91-MSA], pp. 1–48 (1991). <http://hdl.handle.net/1721.1/2356>
7. Challapally, A., Pease, C., Raskar, R., Chari, P.: The GenAI divide: State of AI in business 2025. MIT Project NADA. (2025). https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf
8. KPMG: Global AI pulse Q1 2026. KPMG International. (2026). <https://assets.kpmg.com/content/dam/kpmgsites/sg/pdf/2026/04/global-ai-pulse-q1-2026-report.pdf.coredownload.inline.pdf>
9. Michalski, J., Winters, S., Gunn, D., Holland, J.: AI ROI: The paradox of rising investment and elusive returns. Deloitte. (2025, October 22). <https://www.deloitte.com/global/en/issues/ai/ai-roi-the-paradox-of-rising-investment-and-elusive-returns.html>
10. Mak, R.: Corporate AI sticker shock will force restraint. Reuters. (2026, June 2). <https://www.breakingviews.com/columns/big-view/corporate-ai-sticker-shock-will-force-restraint-2026-06-02>
11. Rosner-Uddin, R.: Amazon scraps AI leaderboard to stop workers chasing usage scores. Financial Times. (2026, May 28). <https://www.ft.com/content/b1a62a7f-6df5-4c90-94ce-64ce9c9961b6?syn-25a6b1a6=1>

12. Bousquette, I.: Why some companies say AI “tokenmaxxing” is key to survival. *Wall Street Journal*. (2026, April 14). <https://www.wsj.com/cio-journal/why-some-companies-say-ai-tokenmaxxing-is-key-to-survival-e699a128>
13. McNeece, M.: Amazon’s AI coding tool deleted a live server and took AWS down for 13 hours. 365i. (2026, February 22). <https://www.365i.co.uk/news/2026/02/22/amazon-kiro-ai-coding-tool-aws-outage>
14. Amazon Staff: Correcting the Financial Times report about AWS, Kiro, and AI. *Amazon News*. (2026, February 20). <https://www.aboutamazon.com/news/aws/aws-service-outage-ai-bot-kiro>
15. Silberling, A.: Meta is having trouble with rogue AI agents. *TechCrunch*. (2026, March 18). <https://techcrunch.com/2026/03/18/meta-is-having-trouble-with-rogue-ai-agents>
16. World Economic Forum: Future of jobs report 2025: 78 million new job opportunities by 2030 but urgent upskilling needed to prepare workforces. *World Economic Forum*. (2025, January 7). <https://www.weforum.org/press/2025/01/future-of-jobs-report-2025-78-million-new-job-opportunities-by-2030-but-urgent-upskilling-needed-to-prepare-workforces>
17. Green, A., Salvi del Pero, A., Verhagen, A.: Artificial intelligence, job quality and inclusiveness: OECD Employment Outlook 2023. *OECD*. (2023, July 11) https://www.oecd.org/en/publications/oecd-employment-outlook-2023_08785bba-en/full-report/artificial-intelligence-job-quality-and-inclusiveness_a713d0ad.html
18. Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., Wilson, N.: The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In: Yamashita, N., Evers, V., Yatani, K., Ding, X., Lee, B., Chetty, M., Touns-Dugas, P. (eds.) *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Assoc. for Computing Machinery, New York 1–22, (2025). <https://doi.org/10.1145/3706598.3713778>
19. Flenady, G., Sparrow, R.: Cut the bullshit: Why GenAI systems are neither collaborators nor tutors. *Teaching in Higher Ed.* (2025), 1–10. <https://doi.org/10.1080/13562517.2025.2497263>
20. Merola, R., Ernst, E., Samaan, D., del Rio-Chanona, M., & Teutloff, O.: The impact of GenAI on jobs, productivity and work organization: A review of the empirical evidence [ILO Brief]. *International Labour Organization*. (2026). <https://doi.org/10.54394/00034342>
21. Mahoney, N., McEntarfer, E., & Wahal, K. What is really happening to jobs? Separating AI hype from reality (Policy Brief). *Stanford Institute for Economic Policy Research (SIEPR)*. (2026, July). <https://siepr.stanford.edu/publications/policy-brief/what-really-happening-jobs-separating-ai-hype-reality>
22. Mintzberg, H.: *Structure in fives: Designing effective organizations*. Prentice Hall, Englewood Cliffs (1993)
23. Perrow, C.: *Normal accidents: Living with high-risk technologies*. Princeton University Press, Princeton (1999)
24. Perrow, C.: A framework for the comparative analysis of organizations. *Amer. Socio. Rev.* **32**(2), 194–208 (1967). <https://doi.org/10.2307/2091811>
25. Mortimore, D.: Moving beyond human-centric organizational designs. *J. Organ. Des.* **13**(2), 65–75 (2024). <https://doi.org/10.1007/s41469-024-00167-z>
26. Carley, K., Gasser, L.: Computational organization theory. In: Weiss, G. (ed.) *Multiagent systems: A modern approach to distributed artificial intelligence*, pp. 299–330. MIT Press, Cambridge (1999)
27. Cyert, R., March, J.: *A behavioral theory of the firm*. 2nd edn. Cambridge University Press, Cambridge (1992)
28. Kahneman, D.: *Thinking, fast and slow*. Farrar, Straus and Giroux, New York (2013)

29. Galbraith, J.: Organization design. Addison-Wesley Publishing Co., Menlo Park (1977)
30. Kuhn, T.: The structure of scientific revolutions. 2nd edn. University of Chicago Press, Chicago (1970)
31. Polyanyi, M.: Personal knowledge: Towards a post-critical philosophy. The University of Chicago Press, Chicago (2015)
32. Lahiri, A.: Inference belief and interpretation in science. A. Lahiri, Kolkata (2020)
33. Parker, S. K., Knight, C.: The SMART model of work design: A higher order structure to help see the wood from the trees. *Human Resource Manag.*, **63**(2), 265–291 (2024). <https://doi.org/10.1002/hrm.22200>
34. Simon, H.: Applying information technology to organization design. *Public Admin. Rev.* **33**(3), 268–278 (1973). <https://doi.org/10.2307/974804>
35. Ocasio, W.: Attention to attention. *Organ. Sci.* **22**(5), 1286–1296 (2011). <https://doi.org/10.1287/orsc.1100.0602>
36. Mortimore, D., Buettner Jr., R. R., Aten, K.: A cognitive slack approach to organizational design. In In: Thomson, R., et al. *Social, Cultural, and Behavioral Modeling. SBP-BRiMS 2024. Lecture Notes in Computer Science*, vol 14972. Springer, Cham. https://doi.org/10.1007/978-3-031-72241-7_11 (2024)
37. Weick, K. E.: Sensemaking in organizations. Sage Publications, Thousand Oaks (1995)
38. March, J. G.: Exploration and exploitation in organizational learning. *Organ. Sci.* **2**(1), 71–87 (1991). <https://doi.org/10.1287/orsc.2.1.71>
39. Cohan, A., Feldman, S., Beltagy, I., Downey, D., Weld, D. S.: Specter: Document-level representation learning using citation-informed transformers. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pp. 2270–2282. Publisher, Location (2020)
40. McInnes, L., Healy, J., Astels, S.: hdbscan: Hierarchical density based clustering. *J. Open Source Software*, **2**(11), (2017). <https://doi.org/10.21105/joss.00205>
41. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *J. Machine Learning Res.* **12**, 2825–2830. (2011).
42. Settles, B.: Active learning literature survey (Computer Sciences Technical Report No. 1648). University of Wisconsin-Madison (2009)
43. Salton, G., & McGill, M. J.: Introduction to modern information retrieval. McGraw-Hill, Columbus (1983)
44. Weick, K. E., Sutcliffe, K. M.: Managing the unexpected. 2nd edn. John Wiley & Sons, San Francisco (2007)