EDITOR: Daniel O'Leary, University of Southern California, USA

DEPARTMENT: AI IN INDUSTRY

Understanding AI through the Lens of Organizational Semiotics

Frank Xing, *National University of Singapore, Singapore*

Keiichi Nakata, *Henley Business School, University of Reading, UK*

Erik Cambria, *Nanyang Technological University, Singapore*

Kecheng Liu, *Henley Business School, University of Reading, UK*

Abstract—AI is gaining new capabilities, creating new expectations, and being deployed in enterprises without a unified organizational understanding. Consequently, a pressing need for managers is to make sense of AI from an organizational perspective, i.e., to understand what AI is and, specifically, how it is different from previous waves of information technology such as group support systems, social media, and/or machine learning. This article sorts the AI-caused challenges reported from different angles. In addition, several organizational semiotic methods are applied to conceptualize a "sign" of AI, providing an anchor for discussing, analyzing, explaining, and evaluating potential issues, as well as strategies to effectively manage these AI-caused challenges.

With the rapid development of large language models (LLMs), large multimodal models, and their deployment to organizations [1], AI's connotation has drastically changed from a humanoid computing machine, a rule-based expert system, to something challenging to understand. Managers may refer to and mean very different things by AI compared to algorithm engineers. This lack of unified understanding is dangerous and may lead to unintended consequences analogous to many not-so-successful information technology (IT) adoptions happened in history. This situation also creates barriers to achieving the goals of harnessing AI's benefits and mitigating its risks within organizations [2].

To address the problem of sense-making across an organization, we apply organizational semiotics, a theory pioneered by Stamper [4] that sees an organization as an information system in which agents employ signs to communicate and perform purposeful actions [3]. The organization's members may have diverse values, beliefs, and intentions, such that coordination is impossible unless the organization's social norms are to some extent "spelt out", using signs and shared interpretation of them. According to organizational semiotics, the social norms are possible contextual behaviors and can be characterized as informal, formal, and technical, as illustrated in the organizational onion (see Figure 1). An organization's norms first develop as the outskirt informal norms, such as culture, beliefs, values, habits and individual behavior.

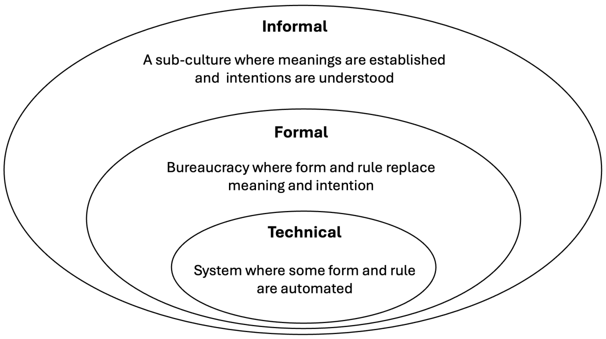

FIGURE 1. The organizational onion, a semiotics view that emphasizes different degrees of social norm formality and institutionalization of signs, adapted from [3].

1541-1672
Digital Object Identifier 10.1109/MIS.2025.XXXXXXX

Some informal norms can gradually be formalized and enter the middle layer as they become more regular, particularly when the efficiency gained through formalization outweighs the risk of signs losing their contextual meaning. This middle formal layer contains bureaucratic forms and rules guiding the individual action, i.e., how work should be done, which can replace meaning and intention in the informal layer with institutionalized signs. As the norms become highly repetitive, IT can be used to automate some parts of the formal systems in the technical layer. In a functional organization, layers provide support for each of the outer layers [5]. Misalignment in containment relationships will have to be adjusted. The organizational onion accords well with Levitt and March's interpretation of organizational learning [6] and helps explain why organizations are eager to adopt new technologies. Since organizational learning is understood as the process of encoding inferences from history into routines that guide behavior — in other words, moving social norms inward — the effective use and expansion of traditional IT serve to accelerate this learning process.

In terms of AI, however, the organizational onion also reveals several potential conflicts/failure-to-support when AI is deployed within an organization. For instance, when the technical truth of AI confronts the organization members' sub-culture, AI may fail to effectively achieve the organization's goals. In extreme cases, members may even work against AI — the phenomenon has been observed in online communities and among knowledge workers when they feel ideas stolen, social interactions destroyed, and reflexivity undermined by AI [7], [8]. An essential semiotic solution lies in effectively communicating the technical truth of AI as a shared sign across the organization, so that form and rule can be established and expectations can be aligned. In the following section, we conceptualize the current technical truth and develop that sign of AI by examining key elements (agents, activities, artifacts) throughout the entire AI pipeline.

CONCEPTULIZATION OF AI

Liu and Li [3] describe organizational semiotics as a radical subjective paradigm in which reality is seen as a social construct. In such a case, the notion of universal truth is irrelevant. Information is only meaningful when it makes sense to its receiver and data is only valuable when it can be understood by its user. In this paradigm, artifacts such as AI models, performances, or software are secondary. Instead, agents and activities are the foci so that it is possible to talk about ownership and responsibility.

Borrowing the form of a Sartrean assertion: **the use of AI precedes AI**. Following this thread, we identify eight AI stakeholder roles that are involved as:

- 1) *policymaker* — who permits, encourages, and incentivize the use of AI
- 2) *founder* — who owns, coordinates, and creates the AI product
- 3) *curator* — who gathers and selects the data needed for building the AI
- 4) *annotator* — who interprets the meaning of data
- 5) *engineer* — who designs how the AI should work
- 6) *evaluator* — who defines, judges how well the AI outcomes are as per desired
- 7) *user* — who prompts, operates the AI
- 8) *client* — who consumes, benefits from the AI.

These stakeholders can be understood as performing different function roles as defined by the organizational morphology [3]. First, policymakers, users, and clients carry out control activities: reinforce the organization's operations, monitor, and reward/penalize other roles. For example, when clients do not like the AI's effects, they put pressures on evaluators to adjust how they act "judge". Second, founders perform communication activities, informing other agents to take actions. Lastly, curators, annotators, engineers, and evaluators perform substantive activities, directly involved in building and maintaining the AI system.

From the perspective of organizational semiotics, it is more useful to conceptualize AI through the unique substantive activities — focusing on what curators, annotators, engineers, and evaluators do. This approach leads to a signifier that emphasizes the activities done pertaining to data, design, and model — similar to the core of what Mihalcea et al. [9] called "research desiderata" but is rooted in an organizational analysis (see Figure 2).

Comparison with Machine Learning

Although the three elements (data, design, and model) are also important to, for example, machine learning, the current AI performs the activities differently. In traditional machine learning, the data is transparent, usually collected after the task is known. In contrast, when using LLMs or large multimodal models, a large proportion of the data (world knowledge) is opaque and collected *ex ante*. This data is topped up with proprietary data such as enterprise prompt libraries or knowledge bases by integration techniques such as retrieval augmented generation (RAG).

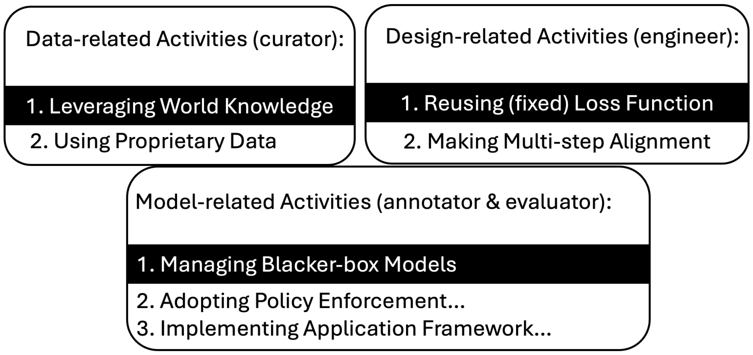

FIGURE 2. A “sign” of the contemporary forms of AI to help organizational understanding. The black backgrounds denotes actitivies external to the target organization.

In terms of design activities, the current AI is less flexible of the loss function. In the first stage, the pure objective is to maximize language coherence. This results in the model not generating the “best” answer but guessing the most “probable” one. In the next stage, some models use a technique called Reinforcement Learning from Human Feedback (RLHF) to align the model behavior to user requirements. This can be done within the organization but is rarely the case because of associated high costs.

These differences render the current AI model a “blacker-box” than traditional machine learning, meaning that not only are the internal mechanics of the model not fully understood, but there is also a loss of partial control over it. This lack of control arises from the inability to locate responsibilities or diagnose problems, especially as some activities occur beyond the organization. Consequently, policies governing when the model can be used and application guidelines for how it should be used must be continuously updated as the base AI model evolves.

SENSE-MAKING OF CHALLENGES WITH ORGANIZATIONAL SEMIOTICS

The sign of AI, combined with the organizational analysis discussed earlier, provides tools for understanding a spectrum of AI challenges we face today. For example, Berente et al. [10] summarized three facets of contemporary forms of AI, i.e., autonomy, learning, and inscrutability. Next, we use the sign of AI to discuss and demonstrate why those facets would emerge, and attempt to answer the questions how they become challenges instead of neutral technological features.

Autonomy, Learning, and Inscrutability

Autonomy refers to the phenomenon that AI has an increasing capacity to act on its own, making decisions such as loan underwriting and staff hiring without human intervention. This phenomenon can be understood as a consequence of design. Previous IT, such as group support systems (GSS), formalized social norms *faithfully*. For example, requiring a booked physical meeting room and a planned agenda for decision-making processes is a highly-shared and repetitive norm. Social media, while influencing how humans interact and controlling the flow of information, still relies on humans to generate contents.

Certain design features of current AI, such as advanced language proficiency, however, enable AI to formalize and perform decision-making tasks, but these tasks often lack the fidelity or depth of human judgment. This can be described as AI’s “Terrible Twos” problem: unlike an infant, it has many capabilities and actions but lacks the maturity and refinement expected of adults. From an organizational perspective, this autonomy poses a challenge because the technical infrastructure (AI, as in Figure 2) fails to make human values explicit or embed them into its design. While human agents tend to delegate their sovereignty to AI to enhance efficiency and productivity, this delegation erodes their ability to take ownership and responsibility for the actions and decisions made by AI.

Learning refers to the phenomenon of AI improving through big data and experience, involving access to data never seen or verified by human agents, thus concerns the “data” element of AI. When AI learns from proprietary data, it raises privacy concerns. When it absorbs large amounts of diverse Internet material to update its world knowledge, it creates distrust. In both cases, learning can be perceived as unsafe by

TABLE 1. Organizational strategies as responses to the AI challenges.

Challenge	Affordance Activities	Organizational Strategies	Example
Autonomy	Mainly design	Permission management	EU AI Act; Risk classification
Learning	Mainly data	Setting up Gen AI guardrails	Explicit quoting; Compliance screening
Inscrutability	Mainly model	Enhance AI literacy	Staff training; Using open-source
(lack of) Coverage	Pan-topical	Production environment dry run	Early client involvement; A/B testing
Biases	Pan-topical	Norm formalization	Enforcing SOP & operation manuals

organization members, blurring roles and responsibilities within an organization.

Inscrutability refers to the phenomenon that AI is only intelligible to some members, but not all members across the organization. This issue arises from differing AI literacy levels and is exacerbated by the evolving nature of policy enforcement and application frameworks, making formalization challenging. As a result, tensions arise among the organization's sub-cultures, hindering organizational learning. Even for the more literate, such as algorithm engineers, black-box models can impose additional costs on substantive activities, as they may need to "reinvent the wheels that are inaccessible from outside the black-box".

Coverage and Biases

Mihalcea et al. [9] illustrated several examples on AI not achieving personal or organizational goals. To summarize, current AI has coverage problems about cultural knowledge and demographic diversity. When things are covered, coverage is often biased. Unlike the previously mentioned three facets, coverage problems and biases are not topical but ubiquitous among agents' activities. For example, biases can be introduced by curators, annotators, or engineers. Consider a healthcare AI system that disproportionately diagnoses mild symptoms as severe conditions without accounting for variations in skin complexion in regions like Africa. The bias could stem from:

- 1) *selection bias* — a curator intentionally gathering less data for darker skin tones.
- 2) *interpretation bias* — an annotator delivering lower-quality work because the same task for darker skin tones are more time-consuming.
- 3) *design bias* — an engineer designing a system that applies lower penalties for diagnostic errors related to darker skin tones.

This example highlights that finer-grained biases can be identified by considering the behaviors of agent roles rather than behaviors of the AI model.

ORGANIZATIONAL STRATEGIES

From an organizational semiotics perspective, AI challenges can be addressed by restoring the containment relationships within the organizational onion. For instance, the autonomy challenge, which arises from the over-extension of technology into less formalized norms, can be mitigated by trimming AI application scopes.

A practical example is the EU AI Act, which classifies Generative AI (Gen AI) systems into four risk categories: "unacceptable risk", "high risk", "limited risk", and "minimal risk". These categories are imposed decreasing compliance pressure and the "unacceptable risk" applications are strictly banned. This approach helps ensure that AI usage is aligned with the organization's formalized structures and values.

To address unexpected and unmonitored learning behaviors, organizations can implement guardrails for Gen AI by explicitly identifying and blocking illegal queries, as well as screening outputs to prevent harmful content or data leaks. Additionally, requiring AI to provide the sources which its responses are based upon can help clarify compliance issues and build trust in its outputs.

Another restoration direction is to expand the scope of both informal and formal norms. This can be achieved by investing in AI literacy training to equip employees with shared knowledge and interpretations, facilitating effective communication and promoting transparency. Establishing norms of inclusiveness and debiasing across all agent roles involved in the AI lifecycle is also essential. For example, when deploying AI for visually impaired individuals, involving them as evaluators can help compare the usefulness of different output forms. Similarly, when introducing digital voice assistants, such as Siri, in multilingual societies like Singapore, engaging local users will uncover how code-switching norms significantly impact the AI's ability to understand languages. These strategies are summarized in Table 1.

SUMMARY

A new way of understanding AI shifts away from viewing it merely as an IT artifact and instead thinks of it as an organizational effort – a collection of activities performed by various stakeholders. This social constructivism approach provides an organizational understanding of the finer-grained structure of current AI and offers insights into the emergence of various AI challenges.

Using organizational semiotics, organizations can develop strategies to harness the benefits of AI while mitigating its associated risks. Most importantly, it allows for accountability by enabling organizations to identify who is responsible when something goes wrong with AI.

REFERENCES

1. D. E. O'Leary, "The rise and design of enterprise large language models," *IEEE Intelligent Systems*, vol. 39, no. 1, 2024, pp. 60–63.
2. E. Cambria, "AI Adoption Phases in Business Intelligence: From Outsourcing to Human-Centered Systems," *Proceedings of IEEE SSCI*, 2025.
3. K. Liu and W. Li, *Organisational semiotics for business informatics*, Routledge, 2015.
4. R. K. Stamper, *Organisational Semiotics: Informatics without the Computer?*, Springer, 2001, pp. 115–171, doi:10.1007/978-1-4615-1655-2_5.
5. A. Jacobs and K. Nakata, "Organisational Semiotics Methods to Assess Organisational Readiness for Internal Use of Social Media," *AMCIS 2012 Proceedings*, 2012, pp. 1–9.
6. B. Levitt and J. G. March, "Organizational Learning," *Annual Review of Sociology*, vol. 14, no. 1, 1988, pp. 319–340.
7. G. Burtch, D. Lee, and Z. Chen, "Generative AI Degrades Online Communities," *Communications of the ACM*, vol. 67, no. 3, 2024, pp. 40–42.
8. D. Lindebaum and P. Fleming, "ChatGPT Undermines Human Reflexivity, Scientific Responsibility and Responsible Management Research," *British Journal of Management*, vol. 35, no. 2, 2024, pp. 566–575.
9. R. Mihalcea et al., "Why AI Is WEIRD and Should Not Be This Way: Towards AI For Everyone, With Everyone, By Everyone," 2024, doi:10.48550/ARXIV.2410.16315.
10. N. Berente et al., "Managing Artificial Intelligence," *MIS Quarterly*, vol. 45, no. 3, 2021, pp. 1433–1450.

Frank Xing is an assistant professor in School of Computing, National University of Singapore. His research focuses on AI for finance, and information systems design. Contact him at xing@nus.edu.sg.

Keiichi Nakata is a professor in Henley Business School, University of Reading. His research interests lie at the interface between technology and people. Contact him at k.nakata@henley.ac.uk.

Erik Cambria is a professor in College of Computing and Data Science, Nanyang Technological University. He is a Fellow of IEEE and his research focuses on neurosymbolic AI and affective computing. Contact him at cambria@ntu.edu.sg

Kecheng Liu is a professor in Henley Business School, University of Reading, and Founding Director of the Informatics Research Centre. He is a BCS Fellow, Fellow of the Charted Institute of Management and an internationally renowned scholar in organizational semiotics. Contact him at k.liu@henley.ac.uk.