

SHODH SAMAGAM


ISSN : 2581-6918 (Online), 2582-1792 (PRINT)



Translating Real-World Problems into Deployable AI/ML Solutions: A Cross-Domain Conceptual Framework and Research Agenda

Naveen Kumar Thawait, Vivek Kumar, Himanshu Yadav, Ajay Sahu
Department of Computer Science & Engineering
Guru Ghasidas Vishwavidyalaya, Koni, Bilaspur, Chhattisgarh, INDIA

ORIGINAL ARTICLE



Authors

**Naveen Kumar Thawait, Vivek Kumar
Himanshu Yadav, Ajay Sahu**

E-mail : thawaitnaveen@gmail.com
shodhsamagam1@gmail.com

Received on : 19/06/2026
Revised on : 20/08/2026
Accepted on : 29/08/2026
Overall Similarity : 00% on 21/08/2026



Plagiarism Checker X - Report

Originality Assessment

0%

Overall Similarity

Date: Aug 21, 2026 (07:39 AM)
Matches: 0 / 3491 words
Sources: 0

Remarks: No similarity found,
your document looks healthy.

Verify Report:
Scan this QR Code



ABSTRACT

Enterprise AI/ML adoption has grown exponentially, yet an increasing share of these initiatives fail to reach production deployment; according to a recent survey, the percentage of abandoned enterprise AI initiative in 2025 is 42%, up from 17% the previous year (S&P Global Market Intelligence, 2025); on the other hand, less than one-third of the surveyed organization uses AI to change core processes substantially (Deloitte AI Institute, 2026). An insightful review of the literature, covering foundations of sociotechnical theory, the healthcare industry, public administration, and industrial AI adoption, shows that there is a common research gap no methodological approach exists for translating ill-structured real-world problems into sociotechnical solutions, which is empirically grounded, cross-domain, and has an ability to generate and deliver successful deployments. The current approaches are domain-specific, not connected to the participatory design paradigm, and lack taxonomy connecting causes in the formulation stage to the reasons of failure in the deployment stage. To propose a conceptual domain-independent framework filling this gap illustrated with three case domains – healthcare, public and industrial/manufacturing. Approach: Critical thematic literature synthesis, prioritization of research gaps according to clear criteria, construction of the conceptual framework, and its roadmap. We put forth the Cross-Domain Problem-to-AI Translation (PAT) framework, which is a seven-layer architecture framework based on the key premise that problem formulation decisions can affect AI systems as much as AI models do. Our gap is discovered independently in four different fields without being imposed upon them, thus fulfilling the condition of real novelty instead of contrived one. We provide a 10–15 year research

roadmap with a clear distinction between the research direction supported by evidence and mere possibilities.

KEY WORDS

AI/ML Problem Formulation, Sociotechnical Systems, Cross-domain Conceptual Framework, AI Deployment Gap, Participatory Design, AI Adoption Barriers.

INTRODUCTION


By 2025, generative and foundation-model capabilities had become widely accessible, and organizational AI adoption accelerated correspondingly: two-thirds of AI-investing firms reported deploying generative AI, while employee access to sanctioned AI tools grew roughly 50% in a single year (S&P Global Market Intelligence, 2025; Deloitte AI Institute, 2026). Yet capability has not translated into commensurate impact, and this gap is neither new nor purely technical. Sociotechnical systems theory established, over seventy years ago, that a technology's effects cannot be understood independently of its social and organizational context (Trist & Bamforth, 1951). Standard process models later formalized steps meant to bridge business need and technical execution most influentially CRISP-DM, whose "business understanding" phase was meant to ground every subsequent decision in the real problem (Wirth & Hipp, 2000). In practice, this phase has remained the least specified stage of the pipeline, even 25 years later. Evidence has accumulated on the consequences. Barocas and Selbst (2016) showed discriminatory outcomes can be encoded during problem definition itself; Obermeyer et al. (2019) demonstrated this empirically at scale in a widely used healthcare algorithm biased by its choice of proxy variable; Passi and Barocas (2019) argued problem formulation is a distinct, undertheorized risk site. Domain-specific frameworks have since emerged independently: in healthcare (El Arab et al., 2025; Jacob et al., 2025), public administration (Madan & Ashok, 2023; Schiff et al., 2021; Babšek et al., 2025), and manufacturing (Gabsi, 2024; Godinho Filho et al., 2025). Each is rigorous within its domain; none engages the others. This fragmentation is itself the paper's central finding: no unified, empirically groundable, cross-domain methodology exists for translating an ill-structured real-world problem into a deployable, sociotechnically accountable AI/ML solution a single gap (Gap A, Section 4) manifesting independently across four literatures rather than four separate gaps, which is precisely what distinguishes it from a manufactured novelty claim.

Existing frameworks fall short for three reasons: no empirical cross-validation exists between domain frameworks; the participatory-design tradition (e.g., Martin et al., 2020) remains disconnected from mainstream ML engineering; and no taxonomy links formulation-stage causes to the independently-named failure modes across literatures. As foundation models commoditize technical capability, the differentiator between AI initiatives that reach production and those abandoned lies increasingly in translation quality, not model quality making a long-term perspective necessary.

This paper is organized around six research questions (RQ1–RQ6, diagnostic through forward-looking) and makes four contributions: a cross-domain literature synthesis; the Cross-Domain Problem-to-AI Translation (PAT) framework, a seven-layer architecture whose central claim is that formulation decisions are as consequential as model architecture; an evidence-graded 10–15 year roadmap; and an explicit statement of validation requirements, making the contribution falsifiable rather than descriptive.

Critical Literature Synthesis

Figure 1: Distribution of reviewed literature by theme (n = 22), reflecting this paper's purposively-compiled matrix (see Section 8 for coverage limitations).



- 1. Process Models and their Blind Spot:** Both process-model traditions reviewed in Section 2 KDD (Fayyad et al., 1996) and CRISP-DM (Wirth & Hipp, 2000) share a common structure: early-stage phases devoted to problem and data understanding, followed by more rigorously specified technical phases. This asymmetry has persisted for nearly three decades without correction, despite both models remaining influential. This persistence suggests the imbalance is a structural feature of how the field has organized its own methodology, not a minor oversight.
- 2. Evidence that Formulation-Stage Choices Cause Downstream Harm:** The fairness-in-ML literature shows a clear progression: Barocas and Selbst (2016) established formulation-stage harm theoretically; Selbst et al. (2019) generalized the mechanism; Obermeyer et al. (2019) confirmed it empirically at scale. None of these papers proposes a validated preventive methodology; Passi and Barocas (2019) frame this explicitly as unresolved.
- 3. Participatory Alternatives and Their Limits:** Martin et al.'s (2020) CBSD operationalizes participatory formulation through collaborative causal modeling, but Sloane et al. (2022) show participation is frequently structured to solicit input without redistributing real authority a tension this paper's framework does not fully resolve (Section 8).
- 4. Domain-Specific Responses: Convergent Diagnosis, Divergent Vocabulary:** Three domain literatures, using different methods and vocabularies, converged on the same diagnosis. Healthcare: El Arab et al. (2025) and Jacob et al. (2025) documented a persistent trial-to-implementation gap. Public administration: Madan and Ashok (2023), Schiff et al. (2021) and Neumann et al. (2022) documented adoption gaps independently; Babšek; et al. (2025) confirmed conceptual fragmentation persists as of 2025. Manufacturing: Gabsi (2024) and Godinho Filho et al. (2025) identified adoption barriers and explicitly flagged the absence of cross-industry comparison. This independent convergence is stronger evidence than any single literature could provide alone.
- 5. The Engineering-Debt Perspective:** Sculley; et al. (2015) found the trained model typically represents a small fraction of a deployed system's complexity and maintenance cost structurally the same argument from an unrelated direction: decisions outside the model are more consequential than the model itself. Amershi; et al. (2019) occupy a similar position for interaction design.

Table 1: Summarizes the literature reviewed above

Author(s), Year	Domain	Core Contribution	Signalled Gap
Trist & Bamforth (1951)	Foundational	Established sociotechnical systems theory: technology and social organization must be jointly optimized	STS theory needs re-theorization for AI-mediated systems
Wirth & Hipp (2000)	Foundational	Proposed CRISP-DM, the dominant process model for data mining/ML projects	“Business understanding” phase remains underspecified relative to technical phases
Barocas & Selbst (2016)	Foundational	Showed discriminatory outcomes can be encoded at the problem-definition stage	No technical framework operationalizes this diagnosis
Passi & Barocas (2019)	Foundational	Named problem formulation as a distinct, undertheorized site of risk	Explicit call for formal methodology, still largely unanswered
Obermeyer et al. (2019)	Healthcare	Empirically demonstrated formulation-stage bias in a deployed healthcare algorithm	Real-world evidence of measurable, deployed harm
El Arab et al. (2025)	Healthcare	Proposed the AI Healthcare Integration Framework (AI-HIF)	Framework untested outside clinical settings
Martin et al. (2020)	Cross-cutting	Proposed Community-Based System Dynamics (CBSD) for participatory formulation	Limited large-scale validation across domains
Sloane et al. (2022)	Cross-cutting	Critiques participation as a procedural “fix” that doesn't redistribute power	Directly complicates CBSD's assumptions — retained as contradictory evidence
Madan & Ashok (2023)	Public sector	Systematic review of AI adoption in public administration	Calls for research bridging strategy and implementation
Schiff et al. (2021)	Public sector	“Public value failure” framework for government AI adoption	Normative framework, not a design methodology
Gabsi (2024)	Industrial	Review of AI integration across Industry 4.0 sectors	Flags absence of cross-sector comparative understanding

Research Gap Analysis

Six candidate gaps were evaluated against eight criteria: scientific importance, novelty, long-term relevance, feasibility, literature support, publication potential, doctoral-research potential, and interdisciplinary reach (Table 2).

Table 2: Summarizes this Evaluation

#	Candidate Gap	Literature Support	Rank
A	No unified, cross-domain, theory-grounded methodology for translating an ill-structured real-world problem into a deployable AI/ML solution	Convergent across foundational (Wirth & Hipp 2000; Passi & Barocas 2019), healthcare (El Arab et al. 2025), public sector (Madan & Ashok 2023), industrial (Gabsi 2024)	central gap
B	Domain frameworks are never empirically cross-validated against one another	Each domain literature stands alone with no comparative study	sub-dimension of A
C	Participatory/stakeholder methods (CBSD, HCAI) and mainstream ML engineering practice (CRISP-DM, MLOps) remain disconnected literatures	Martin et al. 2020 vs. Wirth & Hipp 2000 — essentially no cross-citation between the two traditions	sub-dimension of A
D	No formal failure taxonomy linking problem-formulation root causes to observed deployment failures	“Last-mile,” “sociotechnical gap,” and “public value failure” are named separately with no shared causal model	sub-dimension of A
E	Governance/regulatory frameworks lag sociotechnical transformation, especially Global North–South asymmetries	Selgas-Cors (2025) flags regulatory asymmetry as an open issue	adjacent / future work
F	Standard process models (CRISP-DM) do not formally handle the dynamic complexity/feedback loops that CBSD addresses	Explicit methodological mismatch between the two literatures	sub-dimension of A

The central gap (Gap A) absence of a unified, cross-domain, theory-grounded methodology for translating ill-structured problems into deployable AI/ML solutions is not four separate gaps but one gap independently named by four largely isolated literatures, which is the strongest available evidence of genuineness. Methodological literature named it directly: Wirth and Hipp (2000) left CRISP-DM’s “business understanding” phase unspecified; Passi and Barocas (2019) argued formulation is undertheorized. Domain literatures rediscovered it independently: El Arab; et al. (2025) built AI-HIF because no comparable framework existed; Madan and Ashok (2023) closed their review with the same call; Gabsi (2024) explicitly flagged the cross-sector gap.

Gaps B, C, D, and F resolve into constituent dimensions of Gap A rather than independent problems. Gap E regulatory frameworks lagging AI-driven transformation (Selgas-Cors, 2025) was set aside as adjacent policy scholarship, resurfacing in Section 6’s 2036–2040+ horizon.

The confirmed central gap thus comprises three sub-dimensions: (i) absence of cross-domain empirical validation between existing frameworks; (ii) disconnection between participatory methodologies and mainstream ML engineering; and (iii) absence of a taxonomy linking formulation-stage causes to deployment failure modes.

Proposed Conceptual Framework

Addressing RQ4, this section develops the Cross-Domain Problem-to-AI Translation (PAT) framework (Figure 2). Unlike layered technical architectures common in AI/ML literature (data → learning → explainability → deployment), which describe a trained system’s components, PAT describes stages of translation and validation that precede and surround that system encoding the paper’s central claim that formulation decisions are as consequential to system behavior as architecture decisions.

Figure 2: The Cross-Domain Problem-to-AI Translation (PAT) framework




Table 3: PAT framework layers, function and grounding

Layer	Name	Function	Grounded In / Contribution Type
1	Problem Contextualization	Stakeholder mapping and sociotechnical context scoping	Synthesizes STS theory (Trist & Bamforth 1951; Xu & Gao 2024)
2	Formulation & Specification	Translates contextualized problem into a technical ML specification; formalizes proxy-variable audit steps	Synthesizes CRISP-DM (Wirth & Hipp 2000) with the fairness literature (Barocas & Selbst 2016; Obermeyer et al. 2019)
3	Participatory Validation	Formal checkpoint requiring structured stakeholder review before technical development proceeds	Synthesizes CBSD (Martin et al. 2020); addresses sub-dimension (ii)
4	Technical Development	Standard ML engineering — interfaces with, but does not reinvent, mature MLOps practice	Interfaces with existing literature (Sculley et al. 2015; Treveil et al. 2020); not a novel contribution of this framework
5	Sociotechnical Deployment	Integration into organizational workflows; accountability at go-live	Synthesizes AI-HIF (El Arab et al. 2025) and the public-value-failure model (Schiff et al. 2021)

6	Failure Taxonomy & Monitoring	Classifies post-deployment failures, distinguishing formulation-stage from technical-stage causes	Novel synthesis — addresses sub-dimension (iii)
7	Cross-Domain Feedback	Aggregates formulation records and failure classifications across domains into a shared knowledge base	Novel contribution — addresses sub-dimension (i); closes the central gap

The framework's genuine novelty resides in Layers 6–7, absent from every domain-specific framework reviewed. Layer 6 gives a common vocabulary to phenomena named separately as a “sociotechnical gap” (Ehsan et al., 2023), “public value failure” (Schiff et al., 2021), and unspecified “foundational barriers” (Gabsi, 2024), addressing sub-dimension (iii). Layer 7 aggregates formulation records and failure classifications across domains into a shared knowledge base, addressing sub-dimension (i) and directly targeting the central gap.

Two caveats govern interpretation. First, PAT is a conceptual contribution consistent with the reviewed literature, not an empirically validated methodology. Second, Layers 1–5 synthesize existing, separately-validated components; novelty rests specifically on the integration logic and on Layers 6–7.

10–15 Year Research Roadmap

RQ6 asks how the relative importance of formulation quality is likely to evolve over the next decade and a half. Figure 3 and Table 4 grade four horizons by evidentiary status evidence-supported, plausible projection, or speculative possibility rather than presenting projections as fact.

Figure 3: 10–15 year research roadmap, graded by evidentiary status




Table 4: Roadmap detail: rationale and open questions by horizon

Horizon	Research Direction	Evidentiary Status	Rationale / Current Evidence
2026–2028	Formalizing problem formulation as a distinct engineering discipline within ML pipelines	Evidence-supported	Already underway in fragments: Passi & Barocas (2019) named the gap; AI-HIF and AI for IMPACTS are domain instantiations
2029–2031	Participatory/CBSD-style methods integrated into standard MLOps tooling	Plausible projection	Depends on regulatory pressure and HCAI tooling maturation materializing as expected
2032–2035	AI-assisted problem formulation: foundation models scope/critique formulations before deployment	Plausible projection (higher uncertainty)	Early signals in LLM-assisted requirements engineering; adoption timeline uncertain
2036–2040+	Formal, cross-domain regulatory standards mandating certified translation methodologies	Speculative possibility	Extrapolated from current regulatory trajectory; Selgas-Cors (2025) flags Global North–South regulatory asymmetry

Evidentiary confidence decreases monotonically with time horizon, and each horizon's realization is explicitly conditional on the preceding horizon's research questions being answered, not merely anticipated.

Discussion

This paper's central theoretical claim that formulation quality is co-equal to model architecture implies that evaluation practices privileging benchmark performance may be systematically mis-measuring real-world impact, and testing it requires comparative, cross-domain empirical work treating formulation quality as an independent variable, a genre of study that is currently rare. For practitioners, the most immediately actionable elements are the formulation record and the mandatory participatory checkpoint, both adoptable incrementally within existing workflows.

Limitations

This synthesis draws on 22 purposively selected academic sources rather than an exhaustive, registered systematic search appropriate for a theoretical contribution but narrower than the PRISMA-based reviews it draws on. Sources were identified in English only; given the regulatory asymmetry Selgas-Cors (2025) documents, non-English and Global South scholarship is likely under-represented. Every domain framework reviewed represents a published, successful contribution; failed or abandoned parallel attempts are, by definition, difficult to locate. The domain frameworks compared employ markedly different methods, so their convergence is thematic rather than methodological. The PAT framework itself has not been empirically tested against deployment data, and Sloane et al.'s (2022) participatory-power critique is not resolved by making Layer 3 mandatory. Most fundamentally, this paper proposes a framework; it does not test one.

CONCLUSION

This paper asked whether a scientifically defensible, non-manufactured research gap exists in how AI/ML research addresses translating real-world problems into deployable solutions. The evidence supports an affirmative answer: four largely non-communicating literatures foundational process-model methodology, the fairness-in-ML tradition, and domain-specific research in healthcare, public administration, and manufacturing independently converge on the same diagnosis. In response, this paper proposed the PAT framework, whose Layer 6–7 mechanism directly targets this fragmentation, while remaining explicit that the contribution is conceptual rather than empirical. Researchers building on this work should prioritize, in order: validated formulation-failure taxonomies within single domains, lightweight participatory protocols, and only then the more speculative later-horizon directions. The long-term direction is a field in which formulation quality is measured and audited with the rigor currently reserved for model performance.

BIBLIOGRAPHY

1. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournery, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P. N.; Inkpen, K.; Teevan, J.; Kikin-Gil, R. & Horvitz, E. (2019) Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–13. <https://doi.org/10.1145/3290605.3300233>, Accessed on 04/05/2026.
2. Babšek, M.; Ravšelj, D.; Umek, L. & Aristovnik, A. (2025) Artificial intelligence adoption in public administration: An overview of top-cited articles and practical applications. *AI*, 6(3), 44. <https://doi.org/10.3390/ai6030044>, Accessed on 06/05/2026.
3. Barocas, S. & Selbst, A. D. (2016) Big data's disparate impact. *California Law Review*, 104, 671–732. <https://doi.org/10.2139/ssrn.2477899>
4. Deloitte AI Institute. (2026) The state of AI in the enterprise: The untapped edge (2026 AI report) Deloitte Development LLC, 30 Rockefeller Plaza, New York, NY 10112, USA.
5. Ehsan, U.; Saha, K.; De Choudhury, M. & Riedl, M. O. (2023) Charting the sociotechnical gap in explainable AI: A framework to address the gap in XAI. Proceedings of the ACM on Human-Computer Interaction, 7(CSCW1), Article 34, 1–32. <https://doi.org/10.1145/3579467>, Accessed on 11/05/2026.

6. El Arab, R. A.; Abu-Mahfouz, M. S.; Abuadas, F. H.; Alzghoul, H.; Almari, M.; Ghannam, A. & Seweid, M. M. (2025) Bridging the gap: From AI success in clinical trials to real-world healthcare implementation- A narrative review. *Healthcare*, 13(7), 701. <https://doi.org/10.3390/healthcare13070701>
7. Fayyad, U.; Piatetsky-Shapiro, G. & Smyth, P. (1996) From data mining to knowledge discovery in databases. *AI Magazine*, 17(3), 37–54. <https://doi.org/10.1609/aimag.v17i3.1230>
8. Gabsi, A. E. H. (2024) Integrating artificial intelligence in Industry 4.0: Insights, challenges, and future prospects—A systematic review. *Annals of Operations Research*. <https://doi.org/10.1007/s10479-024-06012-6>, Accessed on 13/05/2026.
9. Godinho Filho, M.; Vieira Queiroz de Almeida, S.; Lage Junior, M.; Osiro, L.; Lima, B. & Callefi, M. H. (2025) A path to follow to overcome foundational barriers to the adoption of artificial intelligence within the manufacturing industry: A conceptual framework. *Enterprise Information Systems*, 19(1–2), Article 2458685. <https://doi.org/10.1080/17517575.2025.2458685>, Accessed on 03/05/2026.
10. Jacob, C.; Brasier, N.; Laurenzi, E.; Heuss, S.; Mougiakakou, S. G.; Cöltekin, A. & Peter, M. K. (2025) AI for IMPACTS framework for evaluating the long-term real-world impacts of AI-powered clinician tools: Systematic review and narrative synthesis. *Journal of Medical Internet Research*, 27, e67485. <https://doi.org/10.2196/67485>, Accessed on 15/05/2026.
11. Madan, R. & Ashok, M. (2023) AI adoption and diffusion in public administration: A systematic literature review and future research agenda. *Government Information Quarterly*, 40(1), 101774. <https://doi.org/10.1016/j.giq.2022.101774>
12. Martin, D., Jr.; Prabhakaran, V.; Kuhlberg, J.; Smart, A. & Isaac, W. S. (2020) Participatory problem formulation for fairer machine learning through community based system dynamics. ICLR 2020 Workshop Paper. arXiv:2005.07572, Accessed on 18/05/2026.
13. Neumann, O.; Guirguis, K. & Steiner, R. (2022) Exploring artificial intelligence adoption in public organizations: A comparative case study. *Public Management Review*. <https://doi.org/10.1080/14719037.2022.2048685>, Accessed on 16/05/2026.
14. Obermeyer, Z.; Powers, B.; Vogeli, C. & Mullainathan, S. (2019) Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
15. Passi, S. & Barocas, S. (2019) Problem formulation and fairness. Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19), 39–48. <https://doi.org/10.1145/3287560.3287567>, Accessed on 20/05/2026.
16. Schiff, D. S.; Schiff, K. J. & Pierson, P. (2021) Assessing public value failure in Government adoption of artificial intelligence. *Public Administration*, 100(2), 427–443. <https://doi.org/10.1111/padm.12742>
17. Sculley, D.; Holt, G.; Golovin, D.; Davydov, E.; Phillips, T.; Ebner, D.; Chaudhary, V.; Young, M.; Crespo, J.-F. & Dennison, D. (2015) Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
18. Selbst, A. D.; Boyd, D.; Friedler, S. A.; Venkatasubramanian, S. & Vertesi, J. (2019) Fairness and abstraction in sociotechnical systems. Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19), 59–68. <https://doi.org/10.1145/3287560.3287598>, Accessed on 22/05/2026.
19. Selgas-Cors, M. (2025) Sociotechnical transformation: A systematic review on the impact of artificial intelligence on society and organizations. *FinTech and Sustainable Innovation*, 1–16. <https://doi.org/10.47852/bonviewFSI52026076>, Accessed on 08/06/2026.

20. Sloane, M.; Moss, E.; Awomolo, O. & Forlano, L. (2022) Participation is not a design fix for machine learning. Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES '22), 177–187. <https://doi.org/10.1145/3551624.3555285>, Accessed on 21/05/2026.
21. S&P Global Market Intelligence. (2025) Generative AI shows rapid growth but yields mixed results (Voice of the Enterprise: AI & Machine Learning, Use Cases 2025), S&P Global Inc., 55 Water Street, New York, NY 10041, USA.
22. Treveil, M.; Omont, N.; Stenac, C.; Lefevre, K.; Phan, D.; Zentici, J.; Lavoillotte, A.; Miyazaki, M. & Heidmann, L. (2020) *Introducing MLOps: How to scale machine learning in the enterprise*. O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472, USA
23. Trist, E. L. & Bamforth, K. W. (1951) Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3–38. <https://doi.org/10.1177/001872675100400101>
24. Wirth, R. & Hipp, J. (2000) CRISP-DM: Towards a standard process model for data mining. Proceedings of the 4th International Conference on Practical Applications of Knowledge Discovery and Data Mining, 29–40. <http://www.crisp-dm.org/>, Accessed on 04/05/2026.
25. Xu, W. & Gao, Z. (2024) An intelligent sociotechnical systems (iSTS) framework: Enabling a hierarchical human-centered AI (hHCAI) approach. arXiv:2401.03223. <https://doi.org/10.48550/arXiv.2401.03223>, Accessed on 06/05/2026.
