International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

Intelligent IT Service Management: Transforming Enterprise Operations Through ServiceNow Automation and Integrated Workflows

Vinod kumar Kolli[1]
[1]*Chicago State University, US*
Email id: *chase9691@gmail.com*[1]

Abstract
Enterprise IT service management (ITSM) is moving from process documentation and manual coordination to data-driven automated service operations. While there are well-defined frameworks that define service-management practices, there is a lack of architectural guidance on how configuration data, artificial intelligence for IT operations (AIOps), predictive models, low-code workflows and digital service channels should work together as one governed system. This conceptual review suggests the five-layer architecture for the intelligent ITSM with the ServiceNow capabilities. The layers include sources and integration, the data foundation, intelligence, orchestration and workflow, and user experience. Governance and continuous improvement is at every level. The framework highlights the importance of the configuration-management database (CMDB), service mapping and transaction history for the reliability of classification, correlation, change-risk scoring and automated fulfilment. It also states that automation doesn't remove human judgement, but rather redistributes it: routine decisions can be automated within limits of confidence, and ambiguous and high consequence cases can be explicitly escalated and held accountable. There are 2 supporting instruments. The first one is a comparison between conventional and intelligent ITSM in the areas of demand intake, routing, knowledge, event management, change evaluation, fulfilment and improvement. The second outlines measurable metrics for service outcomes, automation quality, model performance, and provides cautions for interpretation following process redesign. There is a staged pathway for implementation that starts with data ownership and taxonomy remediation, moves to advisory intelligence and controlled workflow automation, and then evolves to virtual-agent and enterprise-service experiences. The proposed framework provides a foundation for intelligent service operations that is testable, but also needs to be field tested across organizations and platforms.

Keywords: *IT service management; ServiceNow; AIOps; CMDB; workflow automation; predictive intelligence; enterprise service management*

1. Introduction
IT service management is the management of capabilities by which an organization designs, delivers, supports and improves technology-enabled services. The research on ITSM adoption reveals that formal practices might enhance the standardization, transparency and service quality, but the reported gains are diverse as the quality of implementation, organizational commitment and ownership of the process varies significantly [1]-[4]. Another reason for difficulties in benchmarking is the terminological inconsistencies that can occur when organizations refer to the same thing by different names or to different things by the same name [5]. Therefore, the process adoption does not necessarily correspond to operational maturity. The operating environment has also undergone a change. Today, enterprise estates encompass on-premises infrastructure, multiple cloud platforms, software-as-a-service (SaaS) applications, mobile devices and business systems that are constantly evolving. High volumes of events are generated from monitoring tools and service desks are flooded with demand via portals, electronic mail, chat and automated sources. Manual triage and cross functional co-ordination are not on the same growth curve. In response, AIOps research has been born that is centered on anomaly detection, event correlation, failure prediction, root-cause analysis and remediation [6,7]. The opportunity of configurable enterprise platforms. A modern service platform can log work, categorize the demand, correlate technical events, orchestrate fulfilment

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

across departments, suggest knowledge and offer self-service interfaces. But it is an architectural issue to have these capabilities in one package. Inconsistent historical labels cannot be corrected by a classification model, a virtual agent cannot provide useful self-service if fulfilment is still fragmented behind the interface, and event correlation cannot make reliable service impact inferences if there are no component relationships.

This paper proposes an intelligent ITSM simplified reference architecture and takes ServiceNow as the context of the implementation. It answers three questions: what is the proper organization of intelligent-service components to make their dependencies explicit; how is intelligent practice different from traditional process-centred practice; and how can automation be measured without misinterpreting indicators that change their definitions following a redesign. The contribution consists of a layered architecture, a comparative table and a measurement framework. The paper is conceptual and proposes propositions for empirical testing, not claims to demonstrated performance.

2. Literature Background

2.1. ITSM frameworks and organizational embedding

ITSM frameworks are a set of management practices, responsibilities and controls for service delivery. The importance of these is to establish common practices and a foundation for ongoing improvement. However, systematic reviews indicate that the challenges to implementation are primarily of an organisational nature, such as lack of sponsorship, ownership, or adaptation to local context, and resistance to change and excessive process complexity [1, 4]. Other cross-national and empirical studies also attest to the fact that adoption does not create uniform value, but rather is dependent on how practices are integrated in everyday work and how they are in line with business needs [2, 3]. This is important for smart ITSM. A process framework can be consistently categorized, controlled change and measurable service levels, but does not specify whether a judgement is to be done by the requester or by a service agent, by a deterministic rule, or by a learned model. It also does not specify a confidence threshold for automated action, thresholds for evidence to override, or monitoring of model degradation. Only when process control is coupled with computational decision-making do these questions come up.

2.2. Production machine learning and AIOps

AIOps is the use of machine learning and data engineering for operational telemetry and service records. These are anomaly detection, alert reduction, incident prediction, root cause analysis and automated remediation [9]. Mapping studies show that published research is skewed towards detection and fewer studies explain how detected conditions should become part of controlled service workflows [10]. This gap is significant because, for a technically correct anomaly to be operationally valuable, it must have an established pathway to action, evidence, and ownership/priority. Production machine-learning research offers a warning to consider the model as a standalone entity. Instable data-dependencies, undecled consumers, feedback loops and configuration complexity are examples of hidden technical debt [11]. Studies of data lifecycles reveal the need for ongoing engineering effort for collection, validation, feature construction, serving and monitoring. The findings from the software teams in the field also point to the importance of testing, integration, monitoring and coordination for successful deployment, not just model building. These observations support the idea of intelligence being above a governed data foundation and below the controlled execution of workflow.

2.3. Automation of workflows and RPA

Structured inputs, stable rules, and few exceptions are the characteristics of when deterministic automation is suitable. RPA research reports productivity gains for repetitive administrative tasks, and focuses on process selection, exception handling and governance [14]. There are numerous activities that are appropriate for ITSM, including creating accounts, transferring data, routing approvals and fulfilling standard requests. It also encompasses unstructured and ambiguous demand which cannot be safely dealt with by rules alone. Intelligent ITSM is not an either/or proposition between deterministic orchestration and probabilistic inference; it's a combination. The difference will be a factor in control design. A process stage is either finished

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

based on a process rule or an exception is recorded. Rather, a model's output is uncertain and can be correct and yet incorrect. Thus, confidence thresholds, class-specific validations, human review paths and evidence retention are required for model-influenced automation. These controls should be designed prior to delegating decision authority [15].

2.4. Human-AI cooperation and adoption

Intelligent service operations are socio-technical systems. AI management is related to data, model, work design and organizational responsibility management [8]. Human-AI platform research is a hybrid approach where humans and algorithms provide different types of judgements [9]. Evidence from technology-acceptance also reveals that perceived usefulness, effort and social influence are factors that influence the adoption of new channels by employees [16]. Thus, a virtual agent that fails on common requests can have a negative impact on trust, not just in the failed interaction. Automation complements, doesn't replace service jobs. If the routine triage is automated, the rest of the cases are disproportionately ambiguous, novel or consequential. This increases the skill level needed for the first line staff and explanation, escalation and override becomes key part of the way of working. Organizational capability and complementary investments in skills continue to be an important source of technology value [17].

3. Proposed Architecture for Intelligent ITSM

The five horizontal layers in Figure 1 are the governance and continual improvement applied throughout. The architecture is dependency-based - every layer provides reliable services to the layer above it and weakness in one layer limits the performance of the layers above it. This property is the reason why many times failures in the experience or intelligence layer are visible and are caused by failures in integration, data ownership or taxonomy design.

3.1. Sources and integration

The first is the layer that links the service platform to the monitoring system, cloud, identity, human-resources, finance, customer and asset systems. Technical configuration is populated via discovery mechanisms and operational state is exchanged via application interfaces and event streams. The architectural goal is not to be as connected as possible but to have a controlled authority. Each important attribute should have an authoritative source, defined synchronization direction, timestamp and reconciliation rule. The latency requirement will depend on the workflow. Some financial reference data may be able to be updated on a schedule, while employment status for access removal may need to be updated near real-time. Event-driven integration provides better currency, with greater design and observability needs. The integration failures should be reflected as service conditions, as stale or incomplete data from the source can silently lead to wrong decisions in the downstream services [18].

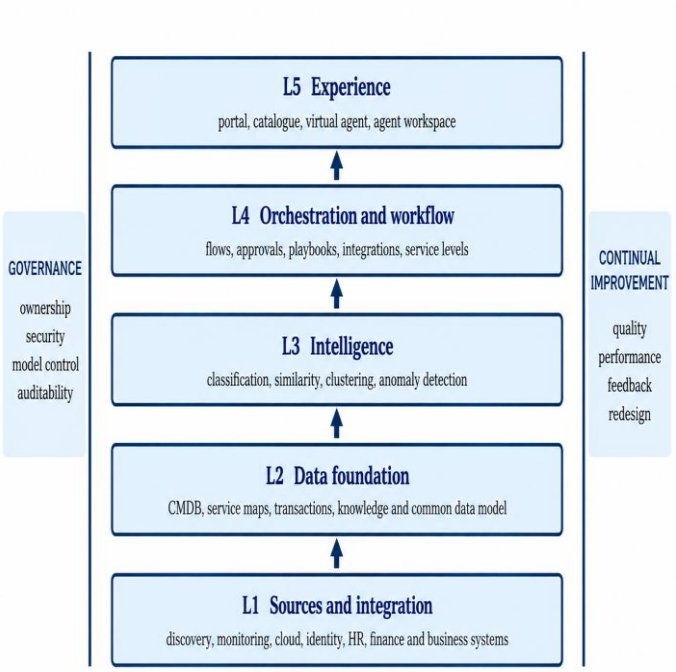

Figure 1 Five-layer architecture for governed intelligent IT service management

3.2. Data foundation

The second layer includes the CMDB, service maps, transactional service records, taxonomies and knowledge. Having a common data model is useful when assets and services have a similar structure, workflows, reports and models can be reused. The CMDB shouldn't be considered a passive inventory. Its relationships give the context to understand how

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

to link a technical component to the business service, owner and users that would be impacted if it failed or changed [19]. This layer is the limiting layer of the architecture. Supervised classification copies systematic labelling mistakes and is based on historical records. Resolution information is incomplete and similarity retrieval is weakened. Event correlation requires topology and service relationships; otherwise, grouping can be done primarily on a time basis, rather than causality. Assessing the risk of change requires an understanding of the impact of a change. Data-quality remediation is thus not a housekeeping task per se, but is part of intelligent-service engineering [20].

3.3. Intelligence

The intelligence layer classifies, retrieves similar data, clusters and detects anomalies and understands language. Classification can recommend category, priority and assignment group. Similarity will bring up related incidents and proven resolutions. Clustering can help determine common needs and issues and gaps in knowledge. Anomaly and event-correlation services can minimize duplicate alerts; and detect patterns that are not part of the normal behaviour. Each inference should be saved along with the version of the model, the score and the outcome. Confidence thresholds should be taken into account of the cost of false action and manual review. A low-risk categorization might be able to accept a different threshold than a high impact change recommendation. Providing an aggregate accuracy is not sufficient as the minority classes may be doing badly and common classes may be doing well. Monitoring of class-specific precision and recall/confusion should be done, especially for security, safety and high priority records [21]. Drift monitoring is critical due to the fact that service catalogues, infrastructure and user language changes. Research on dataset-shift has demonstrated that no detector will work best in all situations [20]. Monitoring should thus include input distributions, prediction confidence, agent corrections, class coverage and operational outcomes. Human corrections should not be taken on faith as rushed and inconsistent corrections can add additional label noise.

3.4. Orchestration and workflow

The fourth layer converts decisions to controlled execution. Low-code flows, approvals, playbooks, integrations and service-level timers orchestrate both automated and human actions. The primary value for the enterprise is likely not to speed up one particular task, but to eliminate coordination delays between departments. For instance, identity, device, payroll, facilities and application-access tasks can be part of a single parent onboarding workflow, with clear ownership and status [22]. It takes software and process discipline to manage workflow artefacts. They should be versioned and tested in controlled environments, reviewed and promoted. Approval authority, segregation of duties or compliance evidence can be modified as much as conventional code in a low code implementation. So automation should be owned, have a recovery behaviour, exception queue and rollback procedure.

3.5. Experience

Portals, catalogues, mobile interfaces, virtual agents and agent workspaces are part of the fifth layer. Guided intake helps to ensure better data quality by collecting structured data up front, before work starts. The initial requests that a virtual agent should support are those that are well-defined, have a high volume, and have high reliability of fulfilment. Broad deployment over weak processes – leads to abandonment and gives users the opportunity to work around the channel. Agent experience is also crucial. Recommendations should be given with supporting evidence, not as "defaults.

Records, knowledge suggestions and likely assignment should decrease search effort whilst providing an obvious means for correction. Explanation methods can help to enhance local understanding of model predictions, but explanations need to be tailored to the user and decision, and not just provided as a technical output [21, 22].

3.6. Governance and continual improvement

Governance covers Data Ownership, Role Based Access, Privacy, Segregation of Duties, Model Approval, Workflow Change Control and Auditability. Ethical-AI lifecycle research highlights the need of assurance throughout the design, deployment, and retirement of AI systems [23]. For every automated decision, there should be an

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

accountable owner, scope and evidence that could be used to reconstruct the reason for the automated decision.

Table 1 Comparison of conventional and intelligent ITSM practice

Process dimension	Conventional practice	Intelligent practice	Principal capability
Demand intake	Telephone, e-mail and free text	Guided portal and conversational intake	Catalogue and virtual agent
Categorization	Manual and inconsistent selection	Model recommendation with confidence threshold	Predictive classification
Assignment	Queue triage and repeated reassignment	Similarity- and skill-informed routing	Similarity and work assignment
Knowledge	Reactive manual search	Contextual recommendation and gap detection	Knowledge and clustering
Events	Separate threshold alerts	Correlation, de-duplication and anomaly detection	AIOps and service mapping
Change	Experience-based committee review	Evidence-based risk and controlled automation	Risk scoring and CMDB relationships
Fulfilment	Departmental hand-offs outside the record	One cross-functional orchestrated workflow	Flows, integrations and playbooks
Improvement	Periodic review after service failure	Continuous signals from outcomes and corrections	Performance analytics and feedback

Continual improvement identifies the next change based on the service outcomes and user feedback, agent corrections and operational incidents. Improvement should be distinguished from structural redesign and local optimization.

Automation of an incorrect process may save time in the short term, but make changes more difficult in the long term. Automation should not be the first step, but rather process simplification and control review.

4. Measurement Framework

Measurement should be a combination of service outcomes, user experience, automation coverage and model quality. Conventional indicators can become redefined following automation. There may be different populations before and after the implementation, as event correlation may decrease the number of incident records. Self-service can create an opportunity to shift some demand from agent-generated records, impacting volume and resolution metrics. Definitions and inclusion rules should thus be pinned down and documented and versioned [22-24]. Baselines should be set prior to redesign. Balancing measures to deter gaming should be used to interpret indicator movement. The first-contact resolution measure should be viewed in conjunction with the reopen rate; the automation coverage measure should be viewed in conjunction with the failure and manual rework rates; and the service-level attainment measure should be viewed in conjunction with the user-experience and outcome measures. Reporting quality of models should be done over time and by class. A brief measurement set and the primary caution associated with each indicator is provided As shown in Table 2 Measurement framework for intelligent IT service management.

5. Implementation Strategy

Start with a bounded service or request family that is known and has volume, pain points and ownership. The organization should capture the current process and identify the desired process and the owners of the service data, process logic and model decisions.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

Table 2 Measurement framework for intelligent IT service management

Indicator	Operational definition	Primary data source	Interpretation caution
Self-service resolution	Confirmed closures without agent involvement / total demand	Interaction, conversation and task records	Exclude abandonment and unconfirmed closure
First-contact resolution	Records resolved during initial contact / resolved records	Incident history and timestamps	Pair with reopen and repeat-contact rates
Mean time to restore	Elapsed time from detected interruption to restored service	Event and incident timestamps	Population changes after correlation can invalidate trends
Reassignment count	Mean assignment-group changes per record	Assignment history	A routing diagnostic, not an independent quality goal
Model coverage and precision	Records above threshold and unchanged predictions	Prediction and correction records	Report by class; aggregates hide minority failure
Automation coverage	Automated fulfilment steps / total eligible steps	Flow execution and task history	Weight by effort and business value
Automation failure rate	Failed or manually recovered executions / all executions	Flow errors and exception queues	Define retries and partial completion consistently
Service-level attainment	Commitments met / commitments due	Service-level records	Separate performance change from target redefinition
Experience score	Standardized post-interaction rating	Survey and experience records	Response bias and channel mix affect comparison
Change failure rate	Changes causing incident, rollback or remediation / implemented changes	Change, incident and problem links	Reliable linking is required to avoid undercounting

Initial measurements should be taken prior to changes. A bounded scope allows to separate the architectural learning from enterprise wide transformation complexity. Prior to the training of supervised models, historical categories, assignment groups and resolution codes should be rationalized. Duplicate and obsolete catalogue items need to be deleted. You should test for completeness the CMDB records and service relationships in the selected scope. Knowledge should be owned, and should have review dates and knowledge links to the services and record types it supports. Often these activities are not as obvious as a new portal or chatbot, but they will be the deciding factor on later reliability [9]. The first model should be in the recommendation mode. Prediction is displayed to agents and agents accept or correct it, no automatic action is taken. This stage delivers realistic data on precision, class imbalance, confidence calibration, user response. The cost of error and review can then be used to select the threshold for future automation. The performance of the model should be measured over time, not over a single test set. The first steps to automate should be those that are stable, have a high volume, and have well-defined exception paths. Technical and process owners, test evidence, permissions review, monitoring and recovery [10] should be provided for each flow. Even if tasks are performed outside of cross-functional workflow, the same status and audit trail should be maintained. It's

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

not about how many automated actions you can do, it's about how you can do them in control. Self-service and conversational channels should be scaled up once fulfilment is consistent. When someone is not confident, interaction design should provide clarity of scope and allow quick transfer to a person. An improvement backlog should be created for agent corrections, failed conversations and workflow exceptions. When expanding to other services, common models, data structures and controls should be reused where appropriate, but do not standardize for the sake of it, because it does not account for significant differences in the domain.

6. Socio-Technical and Governance Considerations

Decision rights need to be clearly defined. Low impact and routine, reversible actions can be automated at a lower evidence level than security, financial or high-availability decisions. A responsibility matrix should be established to determine who is responsible for approving the deployment of the model, establishing thresholds, owns the training data, and accepts operational risk. If human review is to be meaningful, it should not be hundreds of poorly prioritized recommendations that need to be reviewed; that's just moving the automation risk to an overworked person. Explainability should be action oriented. Local explanations can be used to highlight what terms or attributes were important for classification [21] and broader explainable-AI approaches can aid transparency, error analysis and governance [22]. However, an explanation does not constitute evidence that a recommendation is right. Operators must have service context, uncertainty, similar cases, and a well defined way of disagreeing. Causal or policy evidence may be needed for high-stakes recommendations in addition to feature attribution. Monitor user adoption by channel, task and user group. AI service adoption studies indicate that perceived benefit, trust and discomfort are factors that impact the willingness to use AI-enabled services [24]. Completion rate doesn't necessarily reflect that participants are trying multiple times, or dropping out or switching channels. Qualitative feedback and observation can help to determine if an interface decreases effort or simply transfers effort. Levels of authority should also be taken into account in designing automation. The classic automation research differentiates between the acquisition, analysis, selection of information for decision and implementation of action, and suggests adapting the level of automation to reliability and consequence [25]. Intelligent ITSM can do this by automating data collection and suggestions prior to giving permission to execute. Reversible actions with good monitoring can go further than irreversible or highly consequential actions.

7. The risks and control responses

The first risk is that of automating a flawed process. Unneeded approvals and hand-offs can be more efficient, but less apparent. Before implementation, it is important to have process simplification, exception analysis and control review. The second risk is that of silent model degradation. Warning and stop thresholds should be based on monitoring and retraining should be based on evidence rather than a fixed calendar alone [18]. The third risk is lack of confidence in the CMDB. Users find that service relationships are not reliable, which leads to a negative feedback loop and users may stop using the CMDB. Completeness, correctness, compliance and timeliness should be used to assess quality for the services in scope. Remediation should be owned by authoritative sources and process owners and not by one team only [20]. The fourth risk is that of concentration on platforms. The more records, knowledge, workflows and intelligence are consolidated, the more dependency on one operational environment, and the better the coordination. Export paths, integration contracts, recovery and critical workflow logic should be documented in organizations. The fifth risk is access expansion: a single platform can put sensitive service, employee and infrastructure data in play if role design, data segmentation and audit monitoring are not sustained across integrations and analytics created.

Conclusion

Intelligent ITSM can't be accomplished by simply adding a predictive model or conversational interface to an existing service desk. It demands a holistic system where trusted sources deliver to a controlled data foundation, intelligence generates limited

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5978 - 5986
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0766

suggestions, orchestration translates approved decisions into controlled work, and experience helps users and agents. Every layer of governance and continual improvement must be traversed. The main constraint is the CMDB, service maps, taxonomies and transaction history, as the quality of the automation at the higher layers can only be as good as the quality of the automation at the lower layers. The framework also reveals a re-distribution of human judgement in the automation. Routine categorical decisions can be delegated with tested confidence limits, but the ambiguous, novel and high consequence work is still reliant on skilled people. Thresholds, escalation, explanation, accountability and review are thus key design features. The two tables help with implementation by explaining the transition from the traditional practice and by providing measures that are easy to understand even after automation. Future longitudinal and quantitative studies are needed to validate the suggested dependency relationships and the sequence of implementation and its impact on service performance and roles of the workforce.

References

[1]. J. Iden and T. R. Eikebrokk, "Implementing IT service management: A systematic literature review," International Journal of Information Management, vol. 33, no. 3, pp. 512–523, 2013, doi: 10.1016/j.ijinfomgt.2012.11.001.

[2]. M. Marrone and L. M. Kolbe, "Impact of IT service management frameworks on the IT organization," Business & Information Systems Engineering, vol. 3, no. 1, pp. 5–18, 2011, doi: 10.1007/s12599-010-0141-5.

[3]. M. Marrone, F. Gacenga, A. Cater-Steel, and L. Kolbe, "IT service management: A cross-national study of ITIL adoption," Communications of the Association for Information Systems, vol. 34, pp. 865–892, 2014, doi: 10.17705/1CAIS.03449.

[4]. J. Serrano, J. Faustino, D. Adriano, R. Pereira, and M. Mira da Silva, "An IT service management literature review: Challenges, benefits, opportunities and implementation practices," Information, vol. 12, no. 3, Art. no. 111, 2021, doi: 10.3390/info12030111.

[5]. M. Winniford, S. Conger, and L. Erickson-Harris, "Confusion in the ranks: IT service management practice and terminology," Information Systems Management, vol. 26, no. 2, pp. 153–163, 2009, doi: 10.1080/10580530902797532.

[6]. P. Notaro, J. Cardoso, and M. Gerndt, "A systematic mapping study in AIOps," in Service-Oriented Computing—ICSOC 2020 Workshops, Lecture Notes in Computer Science, vol. 12632, 2021, pp. 110–123, doi: 10.1007/978-3-030-76352-7_15.

[7]. Y. Dang, Q. Lin, and P. Huang, "AIOps: Real-world challenges and research innovations," in Proc. IEEE/ACM 41st Int. Conf. Software Engineering: Companion Proceedings, 2019, pp. 4–5, doi: 10.1109/ICSE-Companion.2019.00023.

[8]. N. Berente, B. Gu, J. Recker, and R. Santhanam, "Managing artificial intelligence," MIS Quarterly, vol. 45, no. 3, pp. 1433–1450, 2021, doi: 10.25300/MISQ/2021/16274.

[9]. A. Rai, P. Constantinides, and S. Sarker, "Next-generation digital platforms: Toward human-AI hybrids," MIS Quarterly, vol. 43, no. 1, pp. iii–ix, 2019, doi: 10.25300/MISQ/2019/13304.

[10]. V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," MIS Quarterly, vol. 27, no. 3, pp. 425–478, 2003, doi: 10.2307/30036540.

[11]. A. S. Bharadwaj, "A resource-based perspective on information technology capability and firm performance," MIS Quarterly, vol. 24, no. 1, pp. 169–196, 2000, doi: 10.2307/3250983.

[12]. C. Legner et al., "Digitalization: Opportunity and challenge for the business and information systems engineering community," Business & Information Systems Engineering, vol. 59, no. 4, pp. 301–308, 2017, doi: 10.1007/s12599-017-0484-2.

[13]. S. Aguirre and A. Rodriguez, "Automation of a business process using robotic process automation: A case study," in Applied Computer Sciences in Engineering, vol. 742,

2017, pp. 65–71, doi: 10.1007/978-3-319-66963-2_7.

[14]. R. Syed et al., "Robotic process automation: Contemporary themes and challenges," Computers in Industry, vol. 115, Art. no. 103162, 2020, doi: 10.1016/j.compind.2019.103162.

[15]. D. Sculley et al., "Hidden technical debt in machine learning systems," in Advances in Neural Information Processing Systems, vol. 28, 2015, pp. 2503–2511, doi: 10.48550/arXiv.1504.06520.

[16]. N. Polyzotis, S. Roy, S. E. Whang, and M. Zinkevich, "Data lifecycle challenges in production machine learning: A survey," ACM SIGMOD Record, vol. 47, no. 2, pp. 17–28, 2018, doi: 10.1145/3299887.3299891.

[17]. S. Amershi et al., "Software engineering for machine learning: A case study," in Proc. IEEE/ACM 41st Int. Conf. Software Engineering: Software Engineering in Practice, 2019, pp. 291–300, doi: 10.1109/ICSE-SEIP.2019.00042.

[18]. A. Paleyes, R.-G. Urma, and N. D. Lawrence, "Challenges in deploying machine learning: A survey of case studies," ACM Computing Surveys, vol. 55, no. 6, pp. 1–29, 2023, doi: 10.1145/3533378.

[19]. S. Schelter, D. Lange, P. Schmidt, M. Celikel, F. Biessmann, and A. Grafberger, "Automating large-scale data quality verification," Proceedings of the VLDB Endowment, vol. 11, no. 12, pp. 1781–1794, 2018, doi: 10.14778/3229863.3236237.

[20]. S. Rabanser, S. Günnemann, and Z. C. Lipton, "Failing loudly: An empirical study of methods for detecting dataset shift," in Advances in Neural Information Processing Systems, vol. 32, 2019, doi: 10.48550/arXiv.1810.11953.

[21]. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.

[22]. A. Barredo Arrieta et al., "Explainable artificial intelligence: Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020, doi: 10.1016/j.inffus.2019.12.012.

[23]. R. Ashmore, R. Calinescu, and C. Paterson, "Assuring the machine learning lifecycle: Desiderata, methods, and challenges," Artificial Intelligence Review, vol. 54, pp. 6195–6234, 2021, doi: 10.1007/s10462-021-10093-7.

[24]. D. Gursoy, O. H. Chi, L. Lu, and R. Nunkoo, "Consumers acceptance of artificially intelligent device use in service delivery," International Journal of Information Management, vol. 49, pp. 157–169, 2019, doi: 10.1016/j.ijinfomgt.2019.03.008.

[25]. R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," IEEE Transactions on Systems, Man, and Cybernetics—Part A, vol. 30, no. 3, pp. 286–297, 2000, doi: 10.1109/3468.844354.