

# Advanced Automation for mission-critical Information Technology past AIOps\*

---

Prof. Dr. V. David Sánchez A., Ph.D.  
Brilliant Brains, Palo Alto, California  
November 2024

## Abstract

Best DevOps practices combining software development (Dev) and information technology operations (Ops) shorten the software development lifecycle and support the continuous deployment of high quality software. With increasing requirements to keep the IT infrastructure of enterprises secure (Sec), best practices of DevSecOps are in high demand. Operational (Ops) technologies are in constant development and sometimes definitions might vary from vendor to vendor, but for the most part, there is consensus in terms of what is and what will be needed and time to market of new capabilities is key for an increased business of technology vendors. Along those lines, to be concrete, and without being exhaustive, Figure 1(a) to (d) show how diverse Ops technologies work together, the DataOps dimensions and process, as well as MLOps platforms for data science (DS) and machine learning (ML), respectively. Shortly shy of half of digital initiatives enterprise-wide meets or exceeds their associated business outcome targets according to a recent CIO survey [1]. Outperformers, so called digital vanguards that perform up to twice better than competitors, act differently as the rest strategically as well as tactically. Their growth approach enables partner CxOs to co-lead digital and IT-external efforts to co-build digital capabilities, on the strategic end. Tactically, they take four actions: provide strong foundational platforms, instill architectural awareness, develop skills of business and technology users, and incubate and scale business innovation as summarized in Figure 2. In brief, the most significant reasons for their success is that the digital vanguards co-own the digital delivery and democratize digital capabilities.

For example, the focus of DataOps is the orchestration of people, processes, and technologies to guarantee the quick delivery of high-quality data through, e.g., end-to-end data management and data silos elimination and prevention. Accordingly, there are three DataOps dimensions for its execution: people, processes, and technologies. Utilizing automated processes, quality control, and self-service tools, the DataOps process entails the processing, analysis, learning from, and reusing of data. Five steps are typically executed: pre-collect, collect, organize, analyze, infuse data models as well as monitor quality control. In particular, the infusion of data models into the applications needs to be done in a reproducible and automated fashion. Automated end-to-end monitoring of data pipelines ensures early, less expensive problem identification and remedy. ChatOps can be seen as the utilization/application of chat clients, chatbots, notification systems, chatroom integration tools, and in general real-time comm tools to be integrated into the existing workflow to facilitate, simplify, optimize DevOps tasks based on a conversation-driven collaboration that improves among others information sharing, and feedback loops. Artificial Intelligence for IT Operations (AIOps) increases the IT departments' operational efficiency [3] by combining among others big data and machine learning

---

\*This abstract has been granted permission for public release. The author has been conceiving, designing, building, and operating advanced scalable mission-critical parallel and distributed as well as secure (NSA, commercial) computer systems for federal and state governments and the commercial industry, e.g., AI pipelines deployed to and executing during NASA-ESA-DLR spaceflight missions, ML pipelines for DoD classified programs, advanced AI/ML pipelines for advanced data centers and vertical on-prem, hybrid, and in the cloud applications of the State of California.




Figure 1: Ops technologies (a) Ops cooperation, DevSecOps included in all (b) DataOps dimensions (c) the DataOps process [IBM] (d) MLOps platforms [GARTNER] [2]



Figure 2: Digital Vanguards among CIOs – Agenda 2025: Findings, actions and suggestions [1]

to automate the processes of IT operations, e.g., event correlation, anomaly detection and causality determination. Some initial, fundamental AIOps concepts, architectures, use cases, case studies, applications, and best practices can be found, e.g., in [4, 5]. Monitoring challenges of modern observability of distributed systems with its three pillars: logging, metrics collection, and request tracing have been described, e.g., in [6]. Anomaly detection and alerting, getting the best monitoring data, AI operations and root cause analysis, impact analysis and foundational root causes, auto-remediation, as well as automation and system integrations are described, e.g., in Dynatrace's eBook [7]. Proactively create dashboards, write scripts, manage alerts, and monitor containers using another leading AIOps product Datadog can be inspected, e.g., in [8]. AIOps can have a significant impact in improving key IT Key Performance Indicators (KPIs), including: increasing mean time between failures (MTBF) and decreasing mean time to detect (MTTD), mean time to investigate (MTTI), and mean time to resolution (MTTR) in the five primary use cases for AIOps including: performance analysis, anomaly detection, event correlation and analysis, IT service management, and automation [9]. A self-aware and self-healing machine-driven, automated IT that intelligently helps the ITOps staff to maintain the entire IT estate in the most optimized and efficient manner is called autonomic IT. Key milestones, challenges, and benefits on a journey towards autonomic IT including assessing the organization's progress in five phases: siloed IT monitoring, coordinated IT, machine-assisted IT, AI-advised IT, and autonomic IT is covered in [10]. Examples of currently implemented, planned, and potential AIOps use cases in government agencies like NASA and GSA can be found, e.g., in [11, 12, 13].

To better grasp developments in MLOps and AIOps, a solid background in Artificial Intelligence (AI) and Machine Learning (ML) is indispensable. Some modern foundations of AI/ML can be found in diverse publications including [14, 15], partly showing the racy evolution of those fields from initial treatments, e.g., in [16, 17] after some basic ideas about AI had been introduced in [18, 19]. I myself have been committed to advancing AI and ML since the early 1980's [20] including pioneering breakthrough applications, industrial and in research [21], advanced learning / training methods [22] based on the basic algorithms introduced/described, e.g., in [23, 24, 25], their foundations in and connections to mathematics and statistics as compiled, e.g., in [26, 27] as well as the early use of neurochips [28, 29], advanced development environments [30], and applications of learning control, robotics, and automation as reported in [31, 32, 33]. Reinforcement learning is typically used for control tasks with an intelligent agent attempting to maximize a reward signal while interacting with a dynamic, uncertain environment. Its foundations and advances can be found, e.g., in [34, 35]. Current popular deep learning [36, 37] environments have been compared, e.g., in [38, 39]. Advances in machine learning are continuously being updated, e.g., in [40, 41, 42, 43], with topic intersection in research and applications using big data analytics, cybersecurity, DevSecOps in [44, 45, 46], in computational data science [47], whose foundations, products, and lessons learned from its application in data science projects in academia and industry can be found, e.g., in [48, 49], and key government regulatory issues of AI technology in [50, 51] providing some insight into AI supercomputers for the Gen AI era [52]. Foundations, applications, and recent advances in Generative Artificial Intelligence (Gen AI) and Large Language Models (LLMs) were discussed in [53, 54, 55] and in [56] from a business perspective, just to mention a few of a vast list of publications about this relatively new subject. The provided citations exemplify relevant developments and are by no means meant to be exhaustive.

Between MLOps and AIOps there are sometimes misconceptions regarding their differences and similarities. In principle, they are two completely distinct Ops technologies. The focus of MLOps is the automation and streamlining of the ML modeling lifecycle and the improvement and optimization of their associated development, deployment, maintenance, in general management workflows. The focus of AIOps is the application of AI technologies to automate and optimize IT operations in general, e.g., the maintenance and management of the IT infrastructure including updating the software, IT hygiene, IT security, investigating operational incidents among multiple others. Figure 3 (a) shows the development of a machine learning (ML) solution as an iterative process. At the beginning, a machine learning solution has been identified as an appropriate solution for a given business




Figure 3: MLOps (a) ML as iterative process (b) Large MLOps architecture [AWS]

problem. Then data scientists collect data, integrate data from diverse sources, clean and analyze the data. The process continues by engineering features as well as training, tuning different ML models and evaluating their performance. Based on the result obtained so far, further data collection and additional cleaning might be needed and performed. Model deployment to generate predictions follows if appropriate performance has been achieved. After that step, the ML model is monitored in production. As time passes by in a continuously changing world, the model performance deteriorates to make predictions. Monitoring the model quality allows us to determine when we need to retrain and/or gather new data.

To avoid bottlenecks in the form of manual workflows to upgrade models in production, i.e. to build, train and retrain, deploy, monitor, and manage them, that detrimentally impact the data scientists' productivity, model performance, and costs, we need the appropriate operational practices in place, i.e., an adequate MLOps architecture that fulfill all requirements. The set of operational practices in MLOps supports the automation and standardization of machine learning model building, training, deployment, monitoring, management, and governance. In this fashion, MLOps boosts the data scientists' productivity, maintains a high model accuracy, and enhances security and compliance. For example, the MLOps Orchestrator is a tool that assists in streamlining and enforcing architecture best practices providing an extendable framework to manage ML pipelines for AWS ML and third-party services [57]. Figure 3 (b) shows as an example an AWS MLOps architecture for a staff of 10+ data scientists.

Figure 4 (a) shows the six phases of the ML lifecycle as sequence first: business goal identification, ML problem framing, data processing (data collection, data preprocessing, feature engineering), model development (training, tuning, evaluation), model deployment (inference, prediction), and model monitoring. The data processing is shown as process data and further subdivided in collect data, pre-process data, and engineer features (the latter two denominated prepare data). Figure 4 (b) shows the ML lifecycle phases with feedback loops which reflects the fact that they are in general not sequential in nature. Figure 4 (c) shows the MLOps maturity levels of the incorporation into your enterprise of ML models in production as a journey: initial, repeatable, reliable, and scalable. Emphasis is placed on not incorporating everything at once, but gradually, according to your level of comfort and staff expertise. Figure 4 (d) shows a well-architected ML lifecycle including 6 pillars of best practices for each phase: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability.

Let us now turn our attention to AIOps. There is currently an increased need for business insights improvement using AI/ML, driven by the growing IT system complexity and the ongoing exponential growth of telemetry data. The AIOps vendors' solutions can be technology- or process-centric. Some key features offered include broad data collection and analytical capabilities to support the dynamic organization nature, flexible deployment in dynamic environments, real-time monitoring of diverse technology stacks, and improved contextual and experiential data collection. AIOps customers need




Figure 4: ML lifecycle and MLOps (a) ML lifecycle phases as sequence first (b) with feedback loops (c) MLOps maturity journey (d) well-architected including best practices [58]

to look for broad competency, strong native including autoremediation capabilities, an expansive library of out-of-the-box (OOTB) third-party connectors, insights supply for each entire transaction, from the beginning to the end. The focus needs to remain on the user experience reducing the mean time to resolve (MTTR) and the mean time to identify (MTTI) to enable business growth, e.g., pinpointing service degradation and outages from the perspective of users, not solely from the infrastructure. Figure 5 (a) shows eleven AIOps vendors depending on market presence and current offering strength subdivided according to the Forrester Wave Report [59] in challengers, contenders, strong performers, and leaders. Figure 5 (b) shows the list of the AIOps vendors and their associated products evaluated.

Figure 6 (a) shows the component technologies of every AIOps platform, namely machine learning and big data, combined to automate IT operations processes, including event correlation, anomaly detection, and causality determination. For example, Microsoft Azure's AIOps applies AI/ML technologies for systems, DevOps, and customers to design, build, and operate complex cloud services at scale, as shown in Figure 6 (b). AIOps methodologies allow us to extract insights from data that lead to actions to take. Insights gained serve to detect issues early, diagnose them efficiently, predict to prevent issues, and optimize the associated workflows. Actions derived include mitigate/resolve issues, avert future pain, optimize resource allocation, and improve architecture and process, as shown in Figure 6 (c). The four levels of AIOps maturity are shown in Figure 7 (a): level 0 manual, level 1 dashboard, level 2 recommendations, and level 3 automation, with corresponding incident management: reactive, proactive, predictive, and automated (closed-loop), respectively. Figure 7 (b) shows the four stages of the Azure BRAIN intelligence pipeline: collect, store/process, diagnose/decide, and act. Figure 8 (a) and (b) show the timelines of manual detection of an outage versus a solution using AIOps {SLI (Service Level Indicator) + Auto Detection, Auto Comms, Auto Cross Service Diagnosis, Auto Mitigations} in comparison.

AIOps help competitive organizations to increase efficiency, flexibility, and scale of what used to be their overall DevOps tasks. They need to react fast and adequately to customer demands to secure a predominant spot or enhance their edge in the market. Only a sustained digital transformation via indispensable AIOps can cope with data volume surge, increasing IT infrastructure complexity, decentralized device proliferation, increased expectation towards higher operational efficiency, increased customer and employee satisfaction. In this report, current and future AIOps R&D areas of focused attention are identified and analyzed, in particular innovative AI/ML and GenAI approaches. These include for example, enabling the proactive management of IT environments, optimizing efficiency including downtime minimization by making use of advanced predictive analytics based on the entire breadth and depth of the IT ecosystem data or handling large volumes of real-time, streaming, and telemetry data extending the event processing capabilities of AIOps platforms. New employees



(a)

Vendor	Product evaluated
Datadog	Datadog
Digitate	Ignite AIOps
Dynatrace	Dynatrace
Elastic	Elastic Observability
LogicMonitor	LM Envision
Micro Focus	Operations Bridge
New Relic	Observability Platform
OpsRamp	Platform
ScienceLogic	SL1
Splunk	Observability (Splunk Enterprise/Cloud; Splunk ITSI; Splunk Observability Cloud)
Zenoss	Zenoss Cloud

(b)

Figure 5: Artificial Intelligence for IT Operations (AIOps) (a) AIOps vendors (b) AIOps products [59]




Figure 6: AIOps (a) Definition [Gartner] (b) its AI/ML Technologies (c) its Methodologies [Azure] [60]




Figure 7: Azure BRAIN Network of Intelligence (a) AIOps Maturity Levels [DevOps] (b) its Intelligence Pipeline [60]



Figure 8: Real outage example (a) Manual Detection (b) Solution using AIOps [Azure]

and business employees will be able to acquire institutional knowledge faster, will quickly have a higher impact since the Root-Cause Analysis (RCA) diagnosis can be sped up by GenAI creating customized dashboards and reports. Other GenAI advances can lead to conversational AI chatbots with intrinsic real-time knowledge of the IT ecosystem to enhance in a sophisticated fashion the communication and interaction with users to solve IT problems.

With the goal of reducing or eliminating downtime, AI/ML algorithms are to be leveraged to enhance the RCA capabilities, i.e., to identify and address the root causes of IT issues better. Then, valuable information can be automatically extracted from a large volume of logs, metrics, events, and other sources to enhance the system reliability and the problem resolution speed. Since the security considerations are no longer siloed off the integrated and collaborative environment incorporating DevSecOps and other Ops technologies, a unified approach to information security is to be applied. The development of self-healing AIOps systems with autonomous remediation is being pursued. These new capabilities support guaranteeing strict Service Level Agreements (SLAs) and prevent cybersecurity vulnerabilities or keep them short-lived. AI/ML will make IT management through correlation much more precise and insightful by allowing IT departments to get more value out of the collected IT data. To avoid disconnected silos, public cloud vendor lock-in, higher operational costs, and lacking full visibility across complex, multi-cloud, and hybrid IT environments, AIOps and observability solutions need to provide seamless operations, comprehensive visibility and analytics across diverse infrastructures as opposed to cloud-stack-specific ones. To derive answers from corporate information faster and more holistically, tool silos are to be integrated after the acquisition by conglomerates of vendors in the AIOps, observability, information security, and event management space. All these new capabilities mentioned and their incorporation into advanced AIOps tools and systems of the future are explored and described in detail.

## References

- [1] Gartner, Inc. *2025 CIO Agenda – Unlock superior business outcomes with 4 innovative leadership actions*, 2024.
- [2] Gartner, Inc. *Magic Quadrant for Data Science and Machine Learning*, June 17, 2024.
- [3] MIT Technology Review Insights. *The AIOps mission: Simplify the complex*, MIT Technology Review, August 15, 2019.
- [4] E. Fernandez Climent. *AIOps: Revolutionizing IT Operations with Artificial Intelligence*. ISBN 979-8326046123, May 19, 2024.
- [5] N. Sabharwal and G. Bhardwaj. *Hands-on AIOps – Best Practices Guide to Implementing AIOps*. Apress 2022.
- [6] C. Sridharan. *Distributed Systems Observability – A Guide to Building Robust Systems*. Packt, June 2018.
- [7] Dynatrace. *AIOps done right*. September 22, 2023.
- [8] T.K. Theakanath. *Datadog Cloud Monitoring Quick Start Guide – Proactively create dashboards, write scripts, manage alerts, and monitor containers using Datadog*. Packt, June 2021.
- [9] Splunk. *The Essential Guide to AIOps – Overcome data chaos and get continuous insight into your IT Operations*. 2020.
- [10] ScienceLogic. *Accelerate your Autonomic IT Journey*. 2024.
- [11] GovCIO. *Federal Agencies Building In-House AIOps Expertise*. April 12, 2020  
<https://govciomedia.com/federal-agencies-building-in-house-aiops-expertise/>
- [12] GovCIO. *NASA Focused on Cultural Transformation to Launch AIOps*. June 5, 2021  
<http://govciomedia.com/nasa-focused-on-cultural-transformation-to-launch-aiops/>
- [13] U.S. General Services Administration (GSA). *Emergent Technology – AIOps – Automate IT Operations with AI and ML*. 2024. <https://tech.gsa.gov/emergent-technology/aiops/>
- [14] P. Norvig and S. Russell. *Artificial Intelligence: A Modern Approach, Global Edition*. Pearson, 4th Edition, May 13, 2021.
- [15] V. Smolyakov. *Machine Learning Algorithms in Depth*. Manning, August 27, 2024.
- [16] P.H. Winston. *Artificial Intelligence*. Addison-Wesley, 1st Edition, January 1, 1977.
- [17] M. Minsky and S.A. Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, Expanded Edition, December 28, 1987.
- [18] A.M. Turing. *Computing Machinery and Intelligence*. Mind, 49, 433–460, 1950.
- [19] J. McCarthy et al. *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. August 31, 1955.
- [20] V.D. Sánchez. *Neurocomputing 50th volume anniversary*. Neurocomputing, 50, ix, 2003.
- [21] V.D. Sánchez. – *Neurocomputing - Research and Applications* -; – *Increasing the Autonomy of Space Robots*; – *Intelligent BioSystems*; – *Modeling Dynamics for Communications, Navigation, Guidance and Control Applications*. German Research Center for Anthropotechnics, Wachtenberg-Werthoven, Germany; NASA Ames Research Center, Moffett Field, CA; KAIST Department of BioSystems, Daejeon, South Korea; Rockwell Collins, Advanced Technology Center (ATC), Cedar Rapids, IA, 1990, 2002, 2003, 2011.

[22] V.D. Sánchez. *Neurocomputing – Special Issue on Backpropagation, parts I-IV –, 5[4–6], 6[1–2]; – Special issue on RBF Networks, parts I-II –, 19[1–3], 20[1–3]; Advanced Support Vector Machines and Kernel Methods*, 55[1–2], 1993–1994, 1998, 2003.

[23] D.E. Rumelhart et al. *Parallel Distributed Processing, Volume 1: Explorations in the Microstructure of Cognition: Foundations*. The MIT Press, 1986.

[24] J.L. McClelland et al. *Parallel Distributed Processing, Volume 2: Explorations in the Microstructure of Cognition: Psychological and Biological Models*. The MIT Press, 1986.

[25] V.N. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 2nd Edition, 2000.

[26] M.P. Deisenroth et al. *Mathematics for Machine Learning*. Cambridge University Press 2020.

[27] G. James et al. *An Introduction to Statistical Learning – with Applications in R –*. 2nd Edition, Springer 2023.

[28] AMD. *AMD Instinct MI300X Accelerator – Datasheet*. 2023.

[29] NVIDIA. *NVIDIA H100 NVL GPU – Product Brief*. March 2024.

[30] V.D. Sánchez et al. *Maschinenmarkt – On the Way to Intelligence, Structure and Function of Artificial Neural Networks using Supervised Learning; – The Grey Cells as Example, Analyzed Neural Operations can be realized by VLSI Components (in German) –*, 96[46], 97[3]; *Chip Plus – ANSpec, A Specification Language (in German) –*, 7; *Technische Rundschau – Neurocomputers in Industrial Applications (in German) –*, 82[65]; *Neurocomputing – The Design of a Real-Time Neurocomputer Based on RBF Networks –*, 20, 1990, 1991, 1990, 1998.

[31] V.D. Sánchez et al. *Neural Nets in Robotics (in German)*. *Informationstechnik*, 33 [6], 317–322, 1991.

[32] V.D. Sánchez and G. Hirzinger. *The State of the Art of Robot Learning Control Using Artificial Neural Networks*. in O. Khatib, J.J. Craig, and T. Lozano Perez (Eds.), *The Robotics Review 2*, Cambridge, Massachusetts, The MIT Press, 261–283, 1992.

[33] V.D. Sánchez et al. *Applications of Artificial Neural Networks in Automation, part 5: Neural Networks in the Robot Technology (in German)*. *Automatisierungstechnische Praxis*, 35 [5], 296–305, 1993.

[34] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. Bradford Books, 2nd Edition, November 13, 2018.

[35] M. Lapan. *Deep Reinforcement Learning Hands-On: A practical and easy-to-follow guide to RL from Q-learning and DQNs to PPO and RLHF*. Pack, 3rd Edition, November 12, 2024.

[36] I. Goodfellow et al. *Deep Learning*. The MIT Press 2016.

[37] T.J. Sejnowski. *The Deep Learning Revolution*. The MIT Press 2018.

[38] G. Bosch. *Pytorch vs Tensorflow: A Head-to-Head Comparison*. December 4, 2023.  
<https://viso.ai/deep-learning/pytorch-vs-tensorflow/>

[39] I. Palomares Carrascosa. *Keras vs. JAX: A Comparison*. October 23, 2024.  
<https://www.kdnuggets.com/keras-vs-jax-a-comparison>

[40] M.I. Jordan and T.M. Mitchell. *Machine learning: Trends, perspectives, and prospects*. *Science* 349 [6245], 255–260, 2015.

[41] NVIDIA. *GPU-Accelerated Applications – Catalog*. May 2017.

[42] V.D. Sánchez. *Modern Machine Learning Technology*. December 2017.  
<https://profdrvdsaphd.lima-city.de/documents/ModernMachineLearningTechnology.pdf>

[43] M. Khanuja et ali. *Applied Machine Learning and High-Performance Computing on AWS: Accelerate the development of machine learning applications following architectural best practices*. Pack, December 2022

[44] V.D. Sánchez. *Continuous Big Data Applications in Industry – Modern Development Tools, Distributed Operational and Orchestration Systems, Internet of Things –*. December 2016.  
<https://profdrvdsaphd.lima-city.de/documents/ContinuousBigDataApplications.pdf>

[45] P. Ghavami. *Big Data Analytics Methods – Analytics Techniques in Data Mining, Deep Learning, and Natural Language Processing*. DeGruyter 2020.

[46] R. M. Verma and D.J. Marchette. *Cybersecurity Analytics*. Chapman & Hall, August 29, 2022.

[47] V.D. Sánchez. *Computational Data Science Research and Technology Development – State of the Art –*. January 2023.  
<https://profdrvdsaphd.lima-city.de/documents/ComputationalDataScience.pdf>

[48] J.D. Kelleher and B. Tierney. *Data Science*. The MIT Press, 2018.

[49] M. Braschler et al (Eds). *Applied Data Science – Lessons Learned for the Data-Driven Business*. Springer, 2019.

[50] M. Heikkilä. *The White House just unveiled a new AI Bill of Rights – It's the first big step to hold AI to account*. MIT Technology Review, October 4, 2022.

[51] The White House, Office of Science and Technology Policy (OSTP). *Blueprint for an AI Bill of Rights*. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

[52] V.D. Sánchez. *Deep Impact of Advanced AI Technology Developments – Government Regulatory Measures –*. September 2023.  
<https://profdrvdsaphd.lima-city.de/documents/ArtificialIntelligenceRegulation.pdf>

[53] J. Roberts. *Gaining An Edge In Life & Business With AI: Unleashing the Power of Generative AI and Chat GPT*. Piper Publishing, 2023.

[54] S. Ozdemir. *Quick Start Guide to Large Language Models: Strategies and Best Practices for Using ChatGPT and Other LLMs*. Addison-Wesley 2023.

[55] V.D. Sánchez. *Advanced Gen AI & LLM Foundations and Applications – Paving the way to a more powerful and diverse ML –*. December 2023.  
<https://profdrvdsaphd.lima-city.de/documents/AdvancedGenAILLMs.pdf>

[56] E. Mollick et al. *Generative AI: The Insights You Need from Harvard Business Review*. January 30, 2024.

[57] Amazon Web Services. *MLOps Workload Orchestrator – Implementation Guide*, Release v2.2.2, June, 2024.

[58] Amazon Web Services. *AWS Well-Architected Framework – Machine Learning Lens*, July 5, 2023.

[59] Forrester Research, Inc. *The Forrester Wave Artificial Intelligence For IT Operations, Q4 2022*, December 14, 2022.

[60] J. Sheehan. *Managing Cloud Health with AIOps*, AIOps workshop@ICSE23, May 15, 2023.