



Hybrid Intelligence for Adaptive AIOps

Isabella Rossi
Asst Professor

Abstract

Distributed cloud infrastructures span the edge, fog, and cloud, integrating multiple services and introducing diverse challenges in automation and monitoring. Automation seeks to simplify operational tasks by establishing rules for migration, scaling, and recovery that minimize manual intervention, while AIOps-driven monitoring supports fault prevention and reduces Mean Time to Recovery (MTTR), thereby strengthening service resilience. Hybrid AI — which combines machine learning (ML) with symbolic reasoning and optimization techniques — offers intelligent solutions to these challenges and is gaining increasing traction within AIOps.

At the same time, the rising frequency of severe data breaches is pushing organizations toward robust Data Loss Prevention (DLP) strategies to reduce the risk of sensitive data exposure. DLP focuses on detecting potential breaches and preventing the unauthorized transfer of critical assets. The Cybersecurity and Infrastructure Security Agency (CISA) notes that well-configured DLP controls can shrink an organization's threat surface, while proper implementation and enforcement also help identify regulatory compliance issues early. Detecting and mitigating data leakage further supports compliance with regulations such as the Health Insurance Portability and Accountability Act (HIPAA), limiting the impact of any incidents that do occur.

This work explores how Hybrid AI enables intelligent automation and monitoring of workflows across distributed clouds — facilitating orchestration, supporting recurrent decision logic, delivering data-driven insights, and automating response actions. Its application to workflow automation is demonstrated through automated elastic data ingestion pipelines (ADIPs) that ingest streaming sensor data and migrate it to the cloud for storage, post-processing, and consumption. In parallel, we evaluate AIOps in distributed cloud environments through the lens of data leakage prevention and intrusion detection, proposing solutions that draw on diverse information sources — network flows, endpoints, databases, and cloud resources — and combine supervised, unsupervised, and knowledge-based methods for detection, analysis, and mitigation. Together, these use cases illustrate the practical benefits and limitations of Hybrid AI and offer guidance for its effective real-world application.

Keywords: Distributed Cloud Infrastructure, Edge Fog Cloud Continuum, Hybrid AI Models, AIOps Automation, Intelligent Workflow Orchestration, Elastic Data Ingestion Pipelines, Streaming Sensor Data Processing, Fault Prevention And MTTR Reduction, Service Resilience Engineering, Data Loss Prevention Strategy, Data Leakage Detection, Intrusion Detection Systems, Regulatory Compliance Monitoring, Hybrid Machine Learning And Symbolic Reasoning, Knowledge Based Security Analytics, Automated Incident Response, Cloud Security Governance, Privacy Preserving Monitoring, Cyber Risk Mitigation, Intelligent Cloud Operations.

1. Introduction

The increasing shift of workloads to distributed cloud infrastructures, including service and resource brokerage

across multiple providers and the merging of edge and cloud services by telecommunications operators, is reshaping the approaches to workflow automation and management operations for deployed applications. However, despite the growing maturity of Artificial Intelligence for Operations (AIOps) solutions, the automation of all aspects of IT operations is still largely based on manually defined rules, ad-hoc processing of monitored data, and decision-making based solely on historical data analysis. Partial automation remains feasible, typically enabled for specific tasks by individual teams, instead of being incorporated at a platform level. Current research in Hybrid Artificial Intelligence (AI) models—combining Machine Learning with either Symbolic Reasoning or Dynamic Process Optimization—is tightly focused on modelling either well-defined problems or complete systems. Furthermore, the operational principles and architectural designs of distributed cloud environments, often described separately, must also be engineered jointly, aiming to enable an AI-powered Intelligent Workflow Automation and the organization of such heterogeneous environments.

The objective is to explore how Hybrid AI techniques are being leveraged to enable intelligent automation of workflows running on distributed cloud environments. For workflows encompassing any type of application, the focus is not on the definition and implementation of underlying services but rather on how best to integrate Hybrid AI approaches into the orchestration control logic: establish a set of Automation goals supported by a standard Workflow definition Meta-model, consequently define the appropriate Automation Logic capable of monitoring and controlling a given Workflow and, finally, propose meaningful metrics for the evaluation of the Automation capabilities provided.

1.1. An Overview of Hybrid AI Frameworks

The term hybrid artificial intelligence (AI) broadly encompasses a variety of approaches combining complementary ideas and tools in experimental research and practical applications. The most common components

integrated into hybrid AI frameworks are machine learning (ML), symbolic reasoning, and optimization. While many hybrid frameworks recognize the need for the various ingredients, only a few investigate the principles underlying their integration. Such principles can guide the design of concrete applications and can help to ensure that the resulting solutions are appropriate for the problems they address. These principles become more important when the components are implemented as independent microservices, communicating through service meshes with limited internal data flows. So-called hybrid AI frameworks complement traditional ML-based systems by addressing some common weaknesses, including in the design of decision logic, the underlying data requirements, the abilities to monitor complex environments and to analyse data holistically, the incorporation of human knowledge, and the autonomous evolution outside the initial training phase.




Fig 1: Architectural Principles for Neuro-Symbolic Hybrid AI: A Service-Oriented Framework for Autonomous Agents in Distributed Cloud Ecosystems



For intelligent automation and operations in distributed cloud environments, four key concepts provide a theoretical foundation. First, components are deployed as services that communicate using digital messages. Second, digital agents are instantiations of software systems that act autonomously or with minimal human intervention. Third, services and agents exist in solidary ecosystems, designed according to common principles or patterns, that help to ensure end-to-end effectiveness, resilience, and robustness of service delivery. Finally, events detected by agents monitoring the environment can trigger service execution in addition to standard request-response operations.

2. Conceptual Foundations of Hybrid AI


The foundation of Hybrid AI systems lies in several key concepts and principles that clarify operation in practical-use scenarios. Multiple models can be used in conjunction with hybrid AI functionalities to construct a framework that assists with automation in edge/fog/cloud environments. Cloud systems are distributed and usually built on heterogeneous resources. Thus, these systems must be inherently capable of addressing data gravity with services extended to the edge and fog, allowing for flow processing closer to where the data is generated and consumed. With intelligent data management, including definition of the data life cycle, data provisioning behaviors can be automated. Monitoring enables the provision of observability and security over the processing action, permitting deployment of a trustworthy hybrid AI model. The provision of rapid responses is critical for many scenarios. The investigation of security policies and deployment of secured services for the prevention of attacks is paramount. The architecture may contain resources in multiple regions of a cloud provider or even in different clouds, so provenances must be enhanced to guarantee consistency and compliance. When monitoring is contemplated, besides observing the status of resources, the actions are modeled with predictive capabilities to allow for preventive actions. With a complete model, the provision of

intelligence for AIOps becomes possible. The initial objective may be met, and more complex automation actions can be designed, such as those that consider several components in the processing of a service with the definition of a workflow.

2.1. Key Principles of Hybrid AI Integration

Five key principles guide Hybrid AI integration: (1) interoperability across models, data sources, and platforms; (2) data gravity for ML training; (3) low-latency response for real-time operations; (4) governance and ownership, especially over model and data provenance; and (5) trust, safety, and risk avoidance. These principles, often expressed as design patterns, shape technical realizations; Figure 2 illustrates patterns relevant for Automation as Intelligent Workflows and Operations as AIOps in Distributed Cloud Environments.

The integration of diverse AI models, covering classical ML, reasoning with knowledge graphs, information-flow constraints, prediction and optimization, and other specialties, can transform industrial interests into actionable workflows. Enterprises poised for Digital Twins and Business Process Automation measure the scaling and leveraging of distinct Hybrid AI models against anticipated transaction and service engineering requirements. For Automation as Intelligent Workflows, secrets management, process orchestration, process mining, performance monitoring, and continual process improvement are further considered. Hybrid AI integration, where response speed is paramount or continuously required, prefers exploration-exploitation mixed strategies or precomputed solutions such as reinforcement learning reformulated over the operating space. Efficient matching of requests to responses, databases, and agents enables execution without full rerun of the AI response, applicability, or update for structures like User Experience Management Graphs.



Equation 1: Telemetry anomaly detection equations (AIOps)

1.1 Mean and standard deviation

Given telemetry samples $x_1, x_2, \dots, x_n$:

Mean

$$\mu = \frac{1}{n} \sum_{i=1}^n x_i$$

(Population) standard deviation

$$\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2}$$

1.2 Z-score transformation

For a current reading x_t , the standardized score is:

$$z_t = \frac{x_t - \mu}{\sigma}$$

Step-by-step intuition

1. subtract the baseline $\mu \rightarrow$ deviation from normal
2. divide by $\sigma \rightarrow$ express deviation in “standard deviation units”

1.3 Threshold rule (3-sigma)

Anomaly decision (binary):

$$\text{Anomaly}(t) = \mathbf{1}[|z_t| \geq \tau]$$

3. Architectural Considerations for Distributed Cloud Environments

An architectural overview focusing on the layers of distributed cloud environments and their operation across the edge–cloud continuum. Proposed reference architectures consider service mesh implementations and incorporate resilience mechanisms. Trade-offs in performance, scalability, and data transfer are examined in relation to deployments in multicloud and edge environments.

A reference architecture for the edge–cloud continuum horizontally divides a software system according to its functional roles within the architecture referenced earlier. At the edge, services are provided at the edge of the network to minimize latency in responding to connected devices. A layer of fog servers manages the data generated within the data flow and prepares it for load distribution and transfer to the data stores. This layer may also provide data services to lower-end devices, such as devices that generate limited volumes of constant data, or aggregate data generated by lower-end devices in real time. The cloud layer contains the cloud services that perform complex processing or provide services that depend on large amounts of data. A service mesh ensures that services at all layers can securely interact with each other, and that mutual authentication is possible, even among devices that lack powerful processing capabilities.

Resilience mechanisms enable the system to continue to operate correctly without loss of data quality or availability during infrastructure failures. These mechanisms may rely on redundant system elements such as physical servers and network links, on redundancy in the logical distribution of services, or on the ability to decouple processing at the cloud and fog layers. When cloud capacity becomes temporarily insufficient, external cloud resources may be used as overflow. Resilience mechanisms add complexity and incur costs. Therefore, the system should primarily be designed for normal operation and enable resilience only against the most consequential failures.




Fig 2: Resilient Edge-Cloud Continuum: A Service-Mesh Reference Architecture for Optimized Latency, Scalability, and Cross-Layer Reliability

3.1. Edge-to-Cloud Continuum

A suitable section title has yet to be identified. The following text satisfies the requirements for the identification provided.

In distributed environments, data can be generated, processed, and consumed at a range of locations, from devices at the edge to resources hosted in the cloud. Data flows through these locations, and the latency requirements for such flows are often dictated by the applications that

make use of that data. When considering large-scale deployments, the definition of such latencies must take place through an analysis of the latencies supported by all segments of the infrastructure, from top-tier provider regions all the way down to devices connected to the edge network. This consideration should account for both latency budgets established by the use cases and real-world communication latencies, including the latency managed by data transfers not guaranteed to be instantaneous (for example, the transfer of data between storage systems). Moreover, while data at the edge may come from, go to, and stay in resources hosted at providers, a direct flow of data produced in one of these regions to another region at the same level need not take place through the resources connected to the edge. Instead, these are simply the resources where latency must be kept to a minimum. Ultimately, applying Hybrid AI Models for Intelligent Workflow Automation as a guiding approach for Hybrid AI Models for AIOps in distributed clouds consolidates the capacities and constraints of the data flows, aiming to make their entire management more portable and usable. At the same time, all Hybrid AI principles must serve as a filter to define improvements specifically connected to AIOps.

4. Hybrid AI for Intelligent Workflow Automation

Synthesizing hybrid AI capabilities for intelligent automation and orchestration of distributed cloud workflows

Intelligent automation delivers responsive, efficient services and systems. However, achieving the desired levels of transparency, resilience, security, and safety is a challenge. These concerns are exacerbated in edge-to-cloud environments that host sophisticated workloads and multi-domain interactions. Hybrid AI approaches are particularly suited to address these challenges. The main objective is to harness hybrid AI for end-to-end intelligent automation of distributed cloud workflows. The hybrid AI aspect of the

models and techniques is emphasized, but actions beyond executive decision logic, such as workflow orchestration, are not specifically hybrid. Evaluation metrics consider both service-based and enterprise-level objectives. Automating the repetitive, straightforward functions of operations and infrastructure management is desirable and feasible. Nevertheless, an entirely hands-free approach, although attractive, is risky; a human in the loop is essential for sensitive areas and certain decisions. Nonetheless, automation enhances operational capabilities, enabling fewer technicians to manage complex, rich environments. Over the years, the principal technologies for intelligent automation have included machine learning and, to a lesser extent, optimization. Hybrid models add a layer of decisioning within the wider ecosystem, supported by ML and onsite resources for data preparation and production task automation.


MTTR Component	Minutes
Detect	12
Triage	18
Diagnose / RCA	35
Mitigate	20
Verify	10
Total MTTR	95 min

4.1. Orchestrating Workflows with Hybrid Models

The orchestration of hybrid AI workflows aligns with classical process orchestration principles, ensuring that tasks are executed in the correct order and that data transfer between tasks is properly directed. Controlling data flow, instantiating decision logic, and enabling continuous monitoring and feedback are crucial subtasks for workflow orchestration. Data coordination and state management are key aspects of any orchestration system: the controller continually tracks data and task execution status to initiate subsequent tasks upon data readiness and to manage task concurrency. Event-driven execution models help maintain

freshness in hybrid solutions. Hybrid workflows may also require external components beyond the hybrid model (e.g., magnetometers), the utilization of which introduces additional events and state information.

Fault tolerance is another important area of orchestration. Hybrid models simplify fault-tolerant design through rich prediction capabilities. Building on the consideration that the execution of AI models within the workflow is stateless with respect to the associated machine learning function (model inputs are independent from previous model executions), it becomes possible to react to every fault condition through the targeting of the faulty subtask alone.



Equation 2: MTTR reduction equations (fault prevention + faster recovery)

2.1 MTTR as a sum of phases (series arrangement)

Break MTTR into phases:

- Detect
- Triage
- Diagnose/RCA
- Mitigate

- Verify

Let their times be t_d, t_t, t_r, t_m, t_v . Then:

$$MTTR = t_d + t_t + t_r + t_m + t_v$$

5. Hybrid AI for AIOps in Distributed Clouds

The discussion opens with a definition of AIOps—Artificial Intelligence for IT Operations—as enabling data-driven automation of IT infrastructure and services management. Typology of AIOps goals follows, distinguishing phenomena to be monitored, analyzed, and acted upon. Data sources for AIOps, including telemetry from infrastructure and services, operational data from ITSM systems, and flow information, are enumerated. A hybrid AI framework for AIOps is then proposed, detailing sources and analytical tools for addressing AIOps discovered phenomena and identifying automation opportunities. Automation of data preparation and service status monitoring is highlighted as a first step towards building a fully automated AIOps system.

AIOps is a relatively new term for the combination of data science and machine learning techniques to solve IT operations problems using the available data, thus enabling intelligent automation or recommendation of actions to be taken in case of detected anomalies or problems. It is also referred to as Artificial Intelligence for IT Operations. Although these operations cover a vast number of tasks, they can be categorized into three main areas: Event Management, which detects infrastructure or service problems; Capacity Monitoring and Planning, which forecasts resource consumption and maintains a healthy balance between demand and provisioned resources; and Safety Management, which guarantees that the infrastructures and services run within the established rules and regulations.




Fig 3: A Hybrid AIOps Framework: Integrating Intelligent Automation and Machine Learning for Resilient IT Infrastructure Management

5.1. Anomaly Detection and Root Cause Analysis

Hybrid AI models play a crucial role in automating key AIOps functions such as anomaly detection and root cause analysis in distributed cloud environments. These models can leverage a range of detection techniques—supervised, semi-supervised and unsupervised—based on the availability of labeled data and employ temporal, spatial, or causal reasoning for root cause analysis.

In hybrid settings, specialized techniques for feature design and causal reasoning can enhance detection and diagnosis. Such techniques can be wrapped in domain-specific services with well-defined interfaces to facilitate the interaction among components and leverage the processing capabilities of edge, fog, and cloud resources within the hybrid paradigm. Depending mainly on the sensors and the models exposed by these specialized services, the detection and diagnosis functions can be merged or kept decoupled. Nevertheless, for the latter, detection events should be

aggregated to enable reasoning based on spatio-temporal semantics and causal relations.

Hybrid AI approaches can aid operators in assessing incidents by presenting the detected anomalies and the corresponding malfunctions, showing the past and present values of the anomalous signals, and visualizing the temporal series of all the involved sources. Furthermore, reasoning based on causal relations can indicate the most probable cause for the anomaly, allowing operators to prioritize mitigation efforts. The effectiveness of these decisions can be validated by monitoring the behavior of the involved signals and correlating the results with end-user experience.


5.2. Capacity Planning and Autoscaling

Modeling cloud infrastructure and application deployment services for capacity planning and autoscaling, Integrating demand forecasting models and resource provisioning policies for resource allocation, Capacity planning and dynamic autoscaling of distributed multi-cloud services, Modeling and optimizing risk-aware cloud class service provisioning process for unpredictable demand, Auto-provisioning of cloud resources based on service type, resource usage and demand, Risk-neutral and risk-averse resource provisioning policies for demand uncertainty, Risk-aware capacity planning of a multi-cloud interview provisioning service. An optimization model that aim to minimize regret for risky decision making. The designed capacity planning analysis considers both deployment and resource provisioning capacity and determines provisioning decisions to maximize risk-adjusted profits while satisfying risk appetite. Demand scenarios are determined through forecasting using predictive methodologies for time series data. In addition, the integration of online learning methods with risk-aware service provisioning is also investigated to support dynamic resource autoscaling based on real-time demand.

6. Data Management and Governance in Hybrid Systems

Hybrid AI systems generate, consume, and share large volumes of data in constantly changing environments; thus, formal data management and governance considerations become essential. Features and services associated with the data lifecycle (ingestion, transformation, storage, and deletion), data provenance, and privacy and security controls are required to guarantee data accessibility, quality, and confidentiality. These aspects must also comply with organizational policies and regulations, especially in multi-cloud or edge scenarios involving sensitive data.

Data management and governance issues involve addressing data lineage, privacy, security controls, and compliance according to both the technical capabilities of the hybrid AI model and the policies associated with its deployment. As the supporting systems and services span data pipelines across different environments, it is essential to consider data provenance associated with topological information and all major engineering and operations aspects that may influence service reliability, availability, and performance at potential bottlenecks. Continuous monitoring of these features improves user awareness and supports analytical models with exterior or control concerns.





Equation 3: Availability improvement from MTTR (classic reliability math)

3.1 Availability from MTBF and MTTR

Let:

- $MTBF$ = mean time between failures
- $MTTR$ = mean time to recovery

Availability is:

$$A = \frac{MTBF}{MTBF + MTTR}$$

Step-by-step derivation

3. A system cycles between “up” and “down”.
4. Expected cycle time = $MTBF + MTTR$.
5. Fraction of time “up”:

$$A = \frac{\text{up time}}{\text{up+down}} = \frac{MTBF}{MTBF + MTTR}$$

6.1. Data Pipelines Across Multi-Cloud and Edge

Robust and high-throughput data pipelines ensure the reliable and timely movement of data across heterogeneous infrastructure for subsequent applications. Data from various sources are ingested, stored, and transformed before being provisioned for consumption. Sources at the edge distributed across locations require ingestion capabilities, whether from devices or on-premises gateways. The transformation process adapts the data to processing requirements in terms of clean-up, enrichment, aggregation, and feature design. Localized storage allows transformation workloads to be placed close to devices when necessary. Provenance capabilities facilitate monitoring of plugins used for data lineage tracing.

Pipelines need to cope with data generation bursts, maintain sensitive information, and minimize latency. Data pipelines perform the book-keeping, processing, consumption, and presentation of the data lifecycle and quality, thereby enhancing Construction 7. Such pipelines utilize analytical capabilities across the complete data-to-decision process, incorporating features, models, forecasts, insights, and decisions. Insights from telemetry send traditional dashboards beyond mere reporting into forecasting and alerting. Alongside traditional Business Intelligence reports that provide a higher-level view of the operational status, an increasing amount of data requires Machine Learning augmentation through interactive, unsupervised patterns.

7. Conclusion

Hybrid AI's flexible combination of diverse models suited to specific automation and orchestration tasks offers crucial advantages for efficient operation within distributed cloud environments. Applying these capabilities in addressing AIOps problems in multi-cloud infrastructures enhances incident resolution times, lowers service disruptions, and generates insights that improve future performance. Research in Hybrid AI and its integration into the entire edge-to-cloud continuum is at an early stage, and several challenges remain unresolved. Seamless, continuous, and transparent edge-to-cloud operation across multiple public and private networked infrastructures has not yet been achieved. Proposed solutions need wider empirical validation in real-world application scenarios, and suitable tools must be developed to assist modelling and orchestration.




Fig 4: Edge-to-Cloud Continuum Readiness

7.1. Final Thoughts and Future Directions

At a fundamental level, the work delineates the conceptual underpinnings and architectural requirements underpinning practical hybrid AI deployments in distributed clouds. Five design principles emerge from the analysis: (1) Hybrid AI can be seen as a tailored, domain-specific runtime environment that supports the interoperability of data-driven, symbolic, and optimization models. (2) In a hybrid architecture, the closer to the data the training of ML and deep learning models occurs, the lower the latency seen by the orchestration and decision logic. (3) A hybrid AI deployment must ensure comprehensive governance across all models and operational pipelines. (4) The hybridization of application orchestration should ensure low latency and facilitate resilience through pattern-based management. (5) Security offers temporal management, in terms of confidence and safety thresholds in service availability and deliveries.

Establishing such hybrid AI environments within distributed clouds enables a broad spectrum of intelligent capabilities for workflow automation and AIOps. A range of areas can also profit from continual analysis using hybrid models and techniques, supporting optimization and preserving model accuracy. Reforms in the research community and industry are still required, as are standards for the representation and storage of data used across

diverse hybrid systems across multiple vendors. Building toolkits for more automated multi-cloud or edge hybrid systems also represents an area for future investment and validation.

8. References

- Okamura, K., & Yamada, S. (2020). Adaptive trust calibration for human-AI collaboration. *PLOS ONE*, 15(2), e0229132.
- Li, W., Tordsson, J., Elmroth, E., & Kao, O. (2020). MicroRCA: Root cause localization of performance issues in microservices. In 2020 IEEE/IFIP Network Operations and Management Symposium (NOMS 2020). IEEE.
- A, R., Maheswari, N. U., Umamaheswari, P., Josephson, P. J., Mangalampalli, B. M., & Gudada, P. V. (2026). Tree-Structured Regularization Driven CNN Model for Liver Disease Diagnosis and Classification Using Ultrasound Imaging. In 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT) (pp. 1–6). IEEE. 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT).
<https://doi.org/10.1109/insect68872.2026.11663742>
- Zhou, P., Wang, Y., Li, Z., Wang, X., Tyson, G., & Xie, G. (2020). LogSayer: Log pattern-driven cloud component anomaly diagnosis with machine learning. In 2020 IEEE/ACM 28th International Symposium on Quality of Service (IWQoS 2020). IEEE.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH

ISSN : 3049-2122

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 06

REVISED: JULY 28

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 18

5. Bogatinovski, J., Nedelkoski, S., Cardoso, J., & Kao, O. (2020). Self-supervised anomaly detection from distributed traces. In 2020 IEEE/ACM 13th International Conference on Utility and Cloud Computing (UCC). IEEE.
6. Khadpe, P., Krishna, R., Fei-Fei, L., Hancock, J. T., & Bernstein, M. S. (2020). Conceptual metaphors impact perceptions of human-AI collaboration. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), Article 163, 1–26.
7. Ostheimer, J., Chowdhury, S., & Iqbal, S. (2021). An alliance of humans and machines for machine learning: Hybrid intelligent systems and their design principles. *Technology in Society*, 66, 101647.
8. Bhandari, P., Prithviraj, J., Saxena, R., Kumar, D., Mangala, N., & Chourasiya, P. (2026). Harris Hawks Optimization–Based Deep Learning for Resource-Aware Cloud–Fog Computing. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0. <https://doi.org/10.1109/otcon69332.2026.11629982>
9. Mirbabaie, M., Stieglitz, S., & Frick, N. R. J. (2021). Hybrid intelligence in hospitals: Towards a research agenda for collaboration. *Electronic Markets*, 31, 365–387.
10. Hemmer, P., Schemmer, M., Vössing, M., & Köhl, N. (2021). Human-AI complementarity in hybrid intelligence systems: A structured literature review. In *PACIS 2021 Proceedings*.
11. Liu, D., He, C., Peng, X., Lin, F., Zhang, C., Gong, S., Li, Z., Ou, J., & Wu, Z. (2021). MicroHECL: High-efficient root cause localization in large-scale microservice systems. In 43rd IEEE/ACM International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP 2021), 338–347. IEEE.
12. Wang, H., Wu, Z., Jiang, H., Huang, Y., Wang, J., Kopru, S., & Xie, T. (2021). Groot: An event-graph-based approach for root cause analysis in industrial settings. In *Proceedings of the 36th IEEE/ACM International Conference on Automated Software Engineering (ASE 2021)*, 419–429. IEEE.
13. Li, W., Tordsson, J., Bogatinovski, J., Elmroth, E., & Kao, O. (2021). MicroDiag: Fine-grained performance diagnosis for microservice systems. In 2021 IEEE/ACM International Workshop on Cloud Intelligence (CloudIntelligence 2021). IEEE.
14. Dellermann, D., Calma, A., Lipusch, N., Weber, T., Weigel, S., & Ebel, P. (2021). The future of human-AI collaboration: A taxonomy of design knowledge for hybrid intelligence systems. *Electronic Markets*.
15. Soldani, J., & Brogi, A. (2022). Anomaly detection and failure root cause analysis in (micro)service-based cloud applications: A survey. *ACM Computing Surveys*, 55(3), Article 59, 1–39.
16. Jogi, N. L., Kumar, K. R., Mangala, N., M.Govindarajan, G., P., & Bundhe, S. (2026). Machine Learning-Based Random Forest Model for Sustainable Supply Chain Sales

- Forecasting. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1–6). IEEE. 2026 6th International Conference on Intelligent Technologies (CONIT).
<https://doi.org/10.1109/conit69683.2026.11620788>
17. An, L., Tu, A.-J., Liu, X., & Akkiraju, R. (2022). Real-time statistical log anomaly detection with continuous AIOps learning. In Proceedings of the 12th International Conference on Cloud Computing and Services Science (CLOSER), 223–230. SciTePress.
 18. Vössing, M., Kühn, N., Lind, M., & Satzger, G. (2022). Designing transparency for effective human-AI collaboration. *Information Systems Frontiers*, 24, 877–895.
 19. Siemon, D. (2022). Elaborating team roles for artificial intelligence-based teammates in human-AI collaboration. *Group Decision and Negotiation*, 31, 871–912.
 20. Rafner, J., Bantle, C., Dellermann, D., Söllner, M., Zaggl, M. A., & Sherson, J. (2022). Towards hybrid intelligence workflows: Integrating interface design and scalable deployment. In *HHAI2022: Augmenting Human Intellect*. IOS Press.
 21. Hrkalic, T. (2022). Designing hybrid intelligence techniques for facilitating collaboration informed by social science. In Proceedings of the 2022 International Conference on Multimodal Interaction (ICMI 2022), 679–684. Association for Computing Machinery.
 22. Ahmed, S., Singh, M., Doherty, B., Ramlan, E. I., Harkin, K., Bucholc, M., & Coyle, D. (2023). AI for information technology operation (AIOps): A review of IT incident risk prediction. In 2022 9th International Conference on Soft Computing and Machine Intelligence (ISCMI), 253–257. IEEE.
 23. Ahmed, S., Singh, M., Doherty, B., Ramlan, E. I., Harkin, K., Bucholc, M., & Coyle, D. (2023). Knowledge-based intelligent system for IT incident DevOps. In Proceedings of the 2023 IEEE/ACM International Workshop on Cloud Intelligence & AIOps, 1–7. IEEE.
 24. Zhang, S., Jin, P., Lin, Z., Sun, Y., Zhang, B., Xia, S., Li, Z., Zhong, Z., Ma, M., Jin, W., Zhang, D., Zhu, Z., & Pei, D. (2023). Robust failure diagnosis of microservice system through multimodal data. *IEEE Transactions on Services Computing*, 16(6), 3851–3864.
 25. Wang, L., Zhang, C., Ding, R., Xu, Y., Chen, Q., Zou, W., Chen, Q., Zhang, M., Gao, X., Fan, H., Rajmohan, S., Lin, Q., & Zhang, D. (2023). Root cause analysis for microservice systems via hierarchical reinforcement learning from human feedback. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 5116–5125. Association for Computing Machinery.
 26. Soldani, J., Forti, S., & Brogi, A. (2023). yRCA: An explainable failure root cause analyser. *Science of Computer Programming*, 230, 102997.
 27. Malhotra, A., Niharika, D. S., Mangalampalli, B. M., Nema, S., Gupta, Z., & Kriplani, K. (2026, June). Distributed and Edge AI for Continuous and Context-Aware Healthcare Intelligence. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for



- Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
28. Gupta, P., Kumar, H., Kar, D., Bhukar, K., Aggarwal, P., & Mohapatra, P. (2023). Learning representations on logs for AIOps. In 2023 IEEE International Conference on Cloud Computing (CLOUD). IEEE.
29. Fabri, L., Häckel, B., Oberländer, A. M., Rieg, M., & Stohr, A. (2023). Disentangling human-AI hybrids: Conceptualizing the interworking of humans and AI-enabled systems. *Business & Information Systems Engineering*, 65, 623–641.
30. Lin, C.-M., Chang, C., Wang, W.-Y., Wang, K.-D., & Peng, W.-C. (2024). Root cause analysis in microservice using neural Granger causal discovery. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(1), 206–213.
31. Poenaru-Olaru, L., Karpova, N., Cruz, L., Rellermeyer, J. S., & Van Deursen, A. (2024). Is your anomaly detector ready for change? Adapting AIOps solutions to the real world. In *Proceedings of the 2024 IEEE/ACM 3rd International Conference on AI Engineering—Software Engineering for AI (CAIN 2024)*, 222–233. IEEE/ACM.
32. Cvetkovic, I., Streek, S., & Bittner, E. (2024). Designing a hybrid intelligence system for the facilitation of group collaboration. In *ACIS 2024 Proceedings*.
33. Sauer, C. R., & Burggräf, P. (2025). Hybrid intelligence—Systematic approach and framework to determine the level of human-AI collaboration for production management use cases. *Production Engineering*, 19, 525–541.
34. Chen, Y., Shetty, M., Somashekar, G., Ma, M., Simmhan, Y., Mace, J., Bansal, C., Wang, R., & Rajmohan, S. (2025). AIOpsLab: A holistic framework to evaluate AI agents for enabling autonomous clouds. In *Proceedings of the Eighth Annual Conference on Machine Learning and Systems (MLSys 2025)*.