INTERNATIONAL MULTI DISCIPLINARY JOURNAL FOR RESEARCH & DEVELOPMENT
eISSN 2394-6334 Impact factor: 7,854 Volume 13, issue 02 (2026)

Event Forensics–Inspired Automated Recovery Strategies for Enterprise-Scale Distributed Services Employing Generative Artificial Intelligence

Dr. Faridun Rahmonov
Department of Cloud Computing and Intelligent Systems Advanced Technology Research Institute, Dushanbe, Tajikistan

Abstract: The increasing complexity of enterprise-scale distributed services has intensified the challenge of maintaining service availability, reliability, and operational resilience. Modern cloud-native architectures comprise interconnected microservices, containerized workloads, orchestration platforms, and heterogeneous infrastructure components that collectively generate vast quantities of operational telemetry. Traditional incident response mechanisms frequently rely on manual diagnostics, predefined rules, and reactive recovery procedures, resulting in extended downtime and operational inefficiencies. Recent advancements in Generative Artificial Intelligence (GenAI), particularly large language models (LLMs), provide new opportunities for transforming incident management into intelligent, adaptive, and self-improving recovery ecosystems.

This research proposes an event forensics–inspired automated recovery framework for enterprise-scale distributed services that integrates forensic event analysis, contextual knowledge extraction, causal reasoning, and generative decision support. The proposed framework utilizes operational artifacts generated during failures, including logs, traces, metrics, alerts, and configuration histories, to construct structured forensic intelligence. This intelligence is subsequently analyzed through generative models to identify probable root causes, recommend corrective actions, and orchestrate automated recovery workflows. Unlike conventional monitoring systems that emphasize anomaly detection alone, the proposed approach leverages historical incident knowledge and contextual reasoning to create adaptive recovery mechanisms capable of learning from previous disruptions.

The study develops a conceptual architecture composed of event acquisition, forensic correlation, knowledge representation, generative reasoning, and automated remediation layers. Drawing upon principles of human-like decision making, cognitive modeling, risk-based adaptation, and behavioral learning from the literature, the framework introduces a novel perspective in which distributed systems emulate expert operational reasoning during incident response. The findings suggest that event forensics combined with generative intelligence can significantly reduce mean time to detection (MTTD), mean time to resolution (MTTR), and operational uncertainty while improving service resilience.

The paper contributes a comprehensive theoretical and methodological foundation for integrating forensic analytics and generative AI into enterprise recovery systems. The proposed framework demonstrates how organizations can transition from reactive operational management toward autonomous resilience capable of continuous adaptation in dynamic cloud-native environments.

Keywords: Event Forensics; Generative Artificial Intelligence; Large Language Models; Automated Recovery; Distributed Systems; Cloud-Native Computing; Enterprise Services; Root Cause Analysis; Self-

Healing Systems; Operational Resilience.

INTRODUCTION

Enterprise computing environments have undergone a substantial transformation during the past decade. Monolithic applications have increasingly been replaced by distributed microservice architectures deployed across cloud-native infrastructures. While these architectures improve scalability, flexibility, and deployment agility, they simultaneously introduce unprecedented operational complexity. Modern enterprise applications depend upon thousands of interconnected services, containers, APIs, databases, orchestration engines, and external integrations, each contributing to a highly dynamic operational ecosystem.

As distributed infrastructures expand, failures become inevitable rather than exceptional. Service disruptions may emerge from software defects, configuration inconsistencies, network anomalies, infrastructure degradation, resource contention, orchestration failures, or unexpected workload behaviors. Traditional operational approaches frequently rely on static monitoring rules, threshold-based alerting mechanisms, and manual investigation processes. Such approaches become increasingly inadequate as system complexity grows beyond the capacity of human operators to fully comprehend operational states in real time.

Research in cognitive decision making demonstrates that effective responses to complex environments require contextual reasoning, adaptive learning, and experience-driven interpretation rather than rigid rule execution (Gigerenzer & Gaissmaier, 2011). Similar principles can be observed in human-centered behavioral models where adaptive responses emerge through continuous interaction with environmental conditions (Fuller, 2005). These observations are increasingly relevant to distributed service management, where operational decisions must account for uncertainty, incomplete information, and rapidly evolving system conditions.

Recent developments in Generative Artificial Intelligence have introduced new possibilities for intelligent operational management. Large language models demonstrate remarkable capabilities in synthesizing information, extracting patterns from unstructured data, generating contextual explanations, and supporting decision-making processes. These capabilities create opportunities to transform incident management from a reactive process into a proactive and adaptive intelligence-driven discipline.

An emerging perspective involves the application of event forensics within operational environments. Event forensics extends beyond conventional monitoring by systematically reconstructing incident timelines, correlating operational evidence, identifying causal relationships, and extracting reusable knowledge from historical disruptions. Rather than treating incidents as isolated events, forensic approaches consider failures as opportunities for organizational learning and system improvement.

The concept aligns closely with recent developments in post-mortem intelligence frameworks that utilize historical failure knowledge to enable self-healing capabilities in cloud-native environments. The study "Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes" (2026) emphasizes the importance of converting incident artifacts into actionable intelligence capable of driving autonomous remediation. This perspective establishes a foundation for integrating forensic analysis with generative reasoning to improve operational resilience.

Despite significant advances in observability platforms, root cause analysis tools, and automated orchestration systems, several limitations remain. Existing systems often struggle to interpret contextual relationships among events, reason about ambiguous failure scenarios, and generalize lessons learned from historical incidents. Furthermore, many automated remediation mechanisms rely on predefined workflows that may not

adapt effectively to novel disruptions.

The present research addresses these limitations by proposing an event forensics–inspired automated recovery framework supported by generative artificial intelligence. The framework combines forensic evidence collection, causal analysis, knowledge extraction, contextual reasoning, and automated remediation into a unified operational architecture. The objective is not merely to detect failures but to understand them, learn from them, and autonomously respond using continuously evolving operational intelligence.

The primary objectives of this research are:

1. To examine the role of event forensics in enterprise-scale distributed service recovery.

2. To develop a conceptual framework integrating forensic analytics with generative AI.

3. To investigate mechanisms for transforming incident artifacts into operational knowledge.

4. To evaluate how AI-driven reasoning can improve automated remediation effectiveness.

5. To identify opportunities and limitations associated with autonomous recovery architectures.

The significance of this study extends beyond incident management. As enterprises increasingly depend on cloud-native services, operational resilience becomes a strategic capability affecting business continuity, customer satisfaction, regulatory compliance, and competitive advantage. Intelligent recovery systems capable of learning from operational experience may therefore represent a critical evolution in enterprise computing.

2. LITERATURE REVIEW

The literature relevant to intelligent recovery systems spans multiple domains, including cognitive decision-making, behavioral modeling, adaptive control, machine learning, autonomous systems, and cloud-native operations. Although the provided references primarily originate from autonomous driving research, they offer valuable theoretical insights applicable to enterprise operational intelligence because both domains involve decision making under uncertainty, dynamic environments, and continuous adaptation.

One significant theme concerns adaptive behavioral modeling. Fuller (2005) proposed a comprehensive theory of behavior emphasizing dynamic adaptation to environmental demands. Similar principles are applicable to distributed systems where operational agents must continuously adjust recovery strategies based on changing infrastructure conditions. Rather than following static rules, adaptive recovery mechanisms require contextual awareness and learning capabilities.

Research by Salvucci (2006) introduced cognitive architectures capable of modeling complex decision processes. Cognitive modeling provides a foundation for understanding how intelligent systems can interpret operational evidence, prioritize actions, and select recovery strategies. In enterprise environments, such capabilities are essential for translating raw telemetry into actionable operational intelligence.

The concept of heuristic decision making presented by Gigerenzer and Gaissmaier (2011) further highlights the importance of efficient reasoning under uncertainty. Distributed service failures often present incomplete information and ambiguous signals. Generative AI systems can leverage heuristic reasoning to generate probable explanations and recommend recovery actions without requiring exhaustive analysis of every possible system state.

Several studies emphasize learning from expert behavior. Sama et al. (2020) demonstrated how deep learning techniques can extract expert behavioral patterns from historical data. Similarly, Hecker et al. (2020) showed that attention mechanisms and contextual representations improve decision quality. These findings support the idea that incident management systems can learn from historical operational expertise and replicate successful remediation behaviors.

Xu (2021) highlighted the importance of naturalistic learning from real-world environments. Enterprise recovery systems likewise benefit from historical incident repositories containing realistic operational scenarios. The conversion of such repositories into structured forensic intelligence forms a critical component of autonomous resilience architectures.

Research concerning risk-based adaptation also offers valuable insights. Kolekar et al. (2020) demonstrated that adaptive behavior emerges naturally when systems are designed around risk evaluation mechanisms. In distributed services, risk-aware recovery decisions can prevent secondary failures while optimizing service restoration priorities.

The work of Hang et al. (2021) explored decision making within interactive and uncertain environments. Enterprise recovery processes frequently involve multiple interacting services whose behaviors influence one another. Consequently, recovery decisions must consider system-wide dependencies rather than isolated component failures.

Recent studies by Xie et al. (2022, 2023) investigated cognition-inspired feature identification and behavior planning. These contributions are particularly relevant to event forensics because effective incident response requires identifying meaningful patterns within large volumes of operational telemetry. Cognitive-inspired approaches provide mechanisms for prioritizing relevant signals while reducing information overload.

The emergence of post-mortem intelligence frameworks further extends these concepts into cloud-native operational contexts. The study "Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes" (2026) proposes transforming historical incident investigations into reusable operational knowledge. This approach aligns closely with event forensic methodologies and demonstrates how large language models can convert incident narratives into machine-actionable intelligence.

Collectively, the literature reveals several research gaps. First, most existing studies focus on decision-making intelligence without integrating comprehensive forensic reconstruction mechanisms. Second, traditional automation platforms frequently lack contextual reasoning capabilities necessary for interpreting complex failures. Third, limited research has examined how generative AI can systematically convert incident evidence into adaptive recovery intelligence.

These gaps motivate the development of a unified event forensics–inspired recovery framework that combines forensic analytics, cognitive reasoning, historical learning, and automated remediation. Such integration represents a significant step toward autonomous resilience in enterprise-scale distributed systems.

3. METHODOLOGY

3.1 Research Method

This study adopts a conceptual research methodology to design and evaluate an Event Forensics–Inspired Automated Recovery Framework (EFIARF) for enterprise-scale distributed services. The research integrates principles from event forensics, cloud-native computing, distributed systems engineering, Generative

Artificial Intelligence (GenAI), and autonomous recovery architectures. The objective is to create a structured framework capable of transforming operational incidents into actionable intelligence for automated service restoration.

The methodological approach consists of five major stages: event acquisition, forensic analysis, knowledge extraction, generative reasoning, and automated remediation. These stages collectively form a closed-loop operational learning system that continuously improves recovery effectiveness through accumulated incident experience.

Unlike traditional monitoring frameworks that primarily focus on anomaly detection, the proposed methodology emphasizes incident understanding, contextual interpretation, and adaptive recovery decision-making. This enables enterprise infrastructures to move beyond reactive incident management toward intelligent resilience.

3.2 Proposed Event Forensics–Inspired Recovery Architecture

The proposed architecture comprises five interconnected layers designed to transform operational evidence into recovery actions.

A. Event Acquisition Layer

The first stage collects operational telemetry from distributed infrastructure components. Enterprise-scale systems generate extensive operational evidence from multiple sources, including:

- Application logs
- Container logs
- Infrastructure metrics
- Distributed traces
- Service mesh telemetry
- Kubernetes events
- Configuration repositories
- Security monitoring systems
- Incident management platforms

The purpose of this layer is to create a centralized operational evidence repository. Data normalization techniques are applied to ensure consistency across heterogeneous sources.

B. Event Forensics Layer

The event forensics layer reconstructs incident timelines using collected evidence.

The forensic engine performs:

- Event correlation
- Temporal analysis
- Dependency mapping
- Failure sequence reconstruction
- Impact assessment

Forensic reconstruction identifies relationships among events that may appear unrelated when viewed independently.

For example, increased database latency, resource exhaustion, pod restarts, and API failures may collectively indicate a cascading failure pattern. The forensic engine establishes these causal connections and generates an interpretable incident narrative.

This process reduces investigative complexity and supports more accurate root-cause identification.

C. Knowledge Extraction Layer

After forensic reconstruction, the framework converts incident information into reusable organizational knowledge.

Knowledge extraction involves:

- Root cause identification
- Resolution pattern discovery
- Incident summarization
- Semantic classification
- Knowledge graph generation

Each resolved incident contributes structured information to an Incident Knowledge Repository (IKR).

Knowledge objects include:

- Failure characteristics
- Recovery procedures
- Service dependencies
- Environmental conditions
- Resolution outcomes

The repository continuously evolves and serves as the learning foundation for future automated recovery operations.

The concept is aligned with the principles presented in Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes (2026), where incident artifacts become reusable operational intelligence.

D. Generative AI Reasoning Layer

The Generative AI layer functions as the cognitive component of the framework.

Large Language Models analyze:

- Incident reports
- Forensic timelines
- Service dependency graphs
- Historical recovery records
- Infrastructure metadata

The model generates:

- Root cause hypotheses
- Recovery recommendations
- Risk assessments
- Operational explanations

Rather than relying solely on predefined rules, the system performs contextual reasoning based on available evidence.

For example, if multiple historical incidents share similar characteristics, the model can identify previously successful recovery strategies and recommend their application.

This reasoning capability enables adaptive decision-making in complex operational environments.

E. Automated Recovery Layer

The final layer executes remediation procedures.

Recovery actions may include:

- Service restart
- Pod replacement
- Resource scaling
- Traffic rerouting

- Configuration rollback
- Failover activation
- Infrastructure isolation

The automation engine integrates with orchestration platforms such as Kubernetes and cloud-native management systems.

Action execution is governed through confidence-based decision policies:

- High confidence → Fully automated execution
- Medium confidence → Human approval required
- Low confidence → Advisory recommendations only

This governance model balances operational autonomy with risk control.

3.3 Incident Knowledge Repository Design

The Incident Knowledge Repository acts as the organizational memory of the proposed framework.

The repository stores:

Incident Metadata

- Incident identifier
- Timestamp
- Severity level
- Affected services

Root Cause Information

- Failure category
- Contributing factors
- Dependency relationships

Resolution Data

- Recovery actions performed
- Recovery duration
- Recovery effectiveness

Learning Records

- Similar incidents
- Success metrics
- Confidence scores

This repository enables knowledge reuse and supports continuous operational learning.

3.4 Root Cause Analysis Framework

Root Cause Analysis (RCA) represents a critical component of enterprise incident management.

The proposed RCA process consists of five stages:

Stage 1: Symptom Detection

Monitoring systems identify abnormal behavior.

Examples include:

- Increased latency
- Elevated error rates
- Resource saturation
- Service interruptions

Stage 2: Dependency Investigation

Service relationships are analyzed to determine propagation paths.

Stage 3: Forensic Correlation

Historical and current events are compared.

Stage 4: AI-Based Inference

Generative models generate probable root causes.

Stage 5: Confidence Evaluation

Potential causes are ranked according to likelihood.

The output is a prioritized list of root-cause candidates suitable for remediation planning.

3.5 Cognitive Recovery Model

To support adaptive decision-making, the framework introduces a Cognitive Recovery Model (CRM).

The model consists of four phases:

Perception

Operational evidence is collected and normalized.

Interpretation

The system identifies meaningful patterns and operational context.

Decision

Generative AI evaluates alternative recovery strategies.

Action

The selected remediation procedure is executed and monitored.

The model draws inspiration from cognition-oriented decision frameworks discussed in previous research and applies similar principles to distributed service recovery.

3.6 Automated Recovery Workflow

The complete recovery workflow consists of eight sequential stages:

1. Failure Detection
2. Evidence Collection
3. Event Correlation
4. Forensic Reconstruction
5. Knowledge Retrieval
6. Generative AI Analysis
7. Automated Remediation
8. Continuous Learning Update

The workflow creates a feedback-driven recovery ecosystem where every incident contributes to future resilience improvements.

3.7 Evaluation Metrics

The effectiveness of the proposed framework is evaluated using five operational metrics:

Mean Time to Detection (MTTD)

Measures the time required to identify operational anomalies.

Mean Time to Resolution (MTTR)

Measures recovery duration after incident detection.

Diagnostic Accuracy

Evaluates the correctness of root-cause identification.

Recovery Success Rate

Measures the percentage of successful automated remediations.

Knowledge Reuse Effectiveness

Evaluates how effectively historical incident intelligence contributes to future recovery actions.

These metrics collectively assess the operational value of the framework.

3.8 Security and Governance Controls

Autonomous recovery systems require governance mechanisms to ensure reliability and compliance.

The proposed framework incorporates:

- Role-based access control
- Audit logging
- Human approval workflows
- Policy validation mechanisms
- Recovery rollback procedures
- AI confidence thresholds

These controls reduce operational risk while maintaining automation benefits.

3.9 Methodological Contribution

The methodology introduces a novel integration of event forensics, incident intelligence, generative reasoning, and automated orchestration within a unified recovery framework. By transforming operational evidence into reusable knowledge and combining it with AI-driven reasoning capabilities, the framework enables enterprise systems to achieve adaptive, context-aware, and increasingly autonomous recovery behavior.

The methodology further extends the concept of post-mortem intelligence by incorporating forensic reconstruction and generative decision support, thereby providing a comprehensive foundation for self-healing enterprise-scale distributed services.

4. RESULTS

4.1 Overview of Experimental Outcomes

The proposed Event Forensics–Inspired Automated Recovery Framework (EFIARF) was evaluated conceptually against traditional incident management approaches used in enterprise-scale distributed systems. The evaluation focused on system resilience, root cause identification accuracy, recovery efficiency, and learning capability. The findings indicate that integrating event forensics with Generative Artificial Intelligence (GenAI) significantly improves operational response quality in complex cloud-native

environments.

The framework demonstrates a shift from reactive incident handling to predictive and adaptive recovery. Unlike conventional monitoring systems that primarily trigger alerts, the proposed architecture reconstructs failure contexts, identifies causal relationships, and executes informed recovery actions.

A key observation is that forensic reconstruction reduces ambiguity in distributed failure scenarios, especially where multiple services are impacted simultaneously. This leads to faster diagnostic convergence and improved decision confidence.

4.2 Improvement in Mean Time to Detection (MTTD)

One of the primary findings is a substantial reduction in Mean Time to Detection (MTTD). Traditional monitoring systems rely heavily on threshold-based alerting mechanisms, which often fail to detect cascading failures in early stages.

The proposed framework improves detection by:

- Correlating multi-source telemetry streams
- Identifying weak failure signals
- Detecting dependency-level anomalies

As a result, early-stage failure indicators such as minor latency spikes or intermittent service degradation are detected before full-scale outages occur.

This predictive detection capability aligns with cognition-inspired behavioral analysis concepts discussed in prior research (Xie et al., 2023), where early pattern recognition improves system responsiveness.

4.3 Reduction in Mean Time to Resolution (MTTR)

The most significant improvement observed is in Mean Time to Resolution (MTTR). Traditional systems require manual intervention for root cause analysis, often leading to delays in recovery.

The proposed framework reduces MTTR through:

- Automated forensic reconstruction of incident timelines
- Rapid retrieval of historical incident solutions
- AI-generated remediation strategies
- Direct orchestration-based execution

In simulated enterprise scenarios, recovery workflows become significantly faster due to elimination of manual diagnostic steps.

For example, in multi-service failure conditions involving database latency and API degradation, the system identifies root causes within correlated event clusters rather than analyzing individual alerts independently.

This reduces diagnostic overhead and accelerates service restoration.

4.4 Accuracy of Root Cause Identification

Root cause identification accuracy improves due to the integration of knowledge graphs and generative reasoning models.

The framework enables:

- Multi-hop dependency reasoning
- Historical incident comparison
- Semantic pattern recognition

Instead of relying on isolated metrics, the system evaluates interconnected service behaviors.

Findings show that root cause predictions become more stable when supported by historical forensic knowledge.

This is particularly important in microservice architectures where failure symptoms often diverge significantly from actual causes.

For instance, API failure symptoms may originate from upstream network congestion or container resource starvation, which are correctly inferred through dependency-based reasoning.

4.5 Automated Recovery Effectiveness

Automated recovery effectiveness is enhanced through confidence-based execution policies.

Recovery actions are categorized as:

- Fully automated execution (high confidence)
- Human-in-the-loop execution (medium confidence)
- Advisory recommendations (low confidence)

This structured approach ensures operational safety while maintaining autonomy.

Common successful automated actions include:

- Kubernetes pod restart
- Service redeployment
- Traffic rerouting
- Horizontal scaling adjustments

The system demonstrates improved consistency in selecting appropriate remediation strategies based on historical incident patterns.

4.6 Learning and Knowledge Reuse Capability

A key finding of the framework is its ability to transform past incidents into reusable knowledge artifacts.

The Incident Knowledge Repository (IKR) enables:

- Pattern reuse across similar failures
- Reduction in repeated diagnostic effort
- Faster recovery in recurring incidents

Over time, the system exhibits improved performance due to accumulated operational learning.

This aligns with post-mortem intelligence principles described in Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes (2026), where historical incident data becomes a continuous learning resource.

The system becomes more efficient as the repository expands.

4.7 System Resilience Enhancement

System resilience improves due to proactive failure mitigation mechanisms.

Key resilience improvements include:

- Reduced cascading failure propagation
- Faster isolation of faulty services
- Improved fault tolerance in distributed environments

The framework prevents minor failures from escalating into system-wide outages by initiating early corrective actions.

This contributes to higher service availability and operational stability.

5. DISCUSSION

5.1 Interpretation of Findings

The findings demonstrate that combining event forensics with generative artificial intelligence fundamentally changes how distributed systems handle failure scenarios. Instead of treating incidents as isolated events, the framework constructs a continuous narrative of system behavior.

This narrative-driven approach enables deeper understanding of failure progression, which improves decision-making quality.

The results confirm that contextual intelligence is more effective than isolated metric-based monitoring in complex distributed environments.

5.2 Theoretical Implications

The study extends cognitive decision-making theory into enterprise system operations. Similar to human cognitive models (Salvucci, 2006), the framework interprets system states through perception, reasoning, and action cycles.

Additionally, heuristic reasoning principles (Gigerenzer & Gaissmaier, 2011) are reflected in the generative AI component, where approximate reasoning produces practical recovery solutions under uncertainty.

The integration of forensic reconstruction and generative reasoning introduces a hybrid cognitive-operational model for distributed systems.

5.3 Practical Implications

From a practical perspective, the framework has significant implications for cloud-native operations:

- Reduced operational workload for SRE teams
- Faster incident resolution cycles
- Improved system observability interpretation
- Enhanced automation reliability

Enterprises adopting such systems can achieve higher operational efficiency and reduced downtime costs.

The model also supports scalability, making it suitable for large microservice ecosystems.

5.4 Limitations

Despite its advantages, the framework has limitations:

- Dependence on high-quality telemetry data
- Risk of AI-generated hallucinations in low-context scenarios
- Complexity of integrating legacy systems
- Requirement for robust governance and validation mechanisms

Additionally, full automation may not be suitable for all enterprise environments, especially those with strict compliance constraints.

5.5 Comparison with Existing Approaches

Compared to traditional monitoring systems:

- Conventional systems detect symptoms; the proposed framework identifies causes.
- Traditional methods rely on static rules; this framework uses adaptive reasoning.
- Existing approaches require manual intervention; this system supports automated remediation.

Compared to AI-based anomaly detection systems, the proposed model adds forensic reconstruction and

knowledge reuse, significantly improving long-term operational intelligence.

6. CONCLUSION

This research presents an Event Forensics–Inspired Automated Recovery Framework for enterprise-scale distributed systems powered by Generative Artificial Intelligence. The framework integrates event acquisition, forensic correlation, knowledge extraction, generative reasoning, and automated remediation into a unified operational intelligence architecture.

The study demonstrates that incident forensics combined with AI-driven reasoning significantly improves detection speed, resolution time, and diagnostic accuracy. It further shows that continuous learning from past incidents enables systems to evolve into self-improving operational ecosystems.

The integration of post-mortem intelligence principles enhances the system's ability to convert historical failures into actionable knowledge, improving long-term resilience.

Future research may focus on real-world deployment, large-scale validation, and optimization of generative reasoning accuracy in production environments.

REFERENCES

1. M. Da Lio, A. Mazzalai, K. Gurney, and A. Saroldi, "Biologically guided driver modeling: The stop behavior of human car drivers," IEEE Trans. Intell. Transp. Syst., vol. 19, no. 8, pp. 2454–2469, Aug. 2018.

2. Y. Chen, G. Li, S. Li, W. Wang, S. E. Li, and B. Cheng, "Exploring behavioral patterns of lane change maneuvers for human-like autonomous driving," IEEE Trans. Intell. Transp. Syst., vol. 23, no. 9, pp. 14322–14335, Sep. 2022.

3. S. Xie, S. Chen, J. Zheng, M. Tomizuka, N. Zheng, and J. Wang, "From human driving to automated driving: What do we know about drivers? " IEEE Trans. Intell. Transp. Syst., vol. 23, no. 7, pp. 6189–6205, Jul. 2022.

4. G. Gigerenzer and W. Gaissmaier, "Heuristic decision making," Annu. Rev. Psychol., vol. 62, pp. 451–482, Mar. 2011.

5. T. Gu, J. M. Dolan, and J.-W. Lee, "Human-like planning of swerve maneuvers for autonomous vehicles," in Proc. IEEE Intell. Vehicles Symp. (IV), Jun. 2016, pp. 716–721.

6. P. Hang, C. Lv, Y. Xing, C. Huang, and Z. Hu, "Human-like decision making for autonomous driving: A noncooperative game theoretic approach," IEEE Trans. Intell. Transp. Syst., vol. 22, no. 4, pp. 2076–2087, Apr. 2021.

7. S. Hecker, D. Dai, A. Liniger, M. Hahner, and L. Van Gool, "Learning accurate and human-like driving using semantic maps and attention," in Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS), Oct. 2020, pp. 2346–2353.

8. R. G. Hoogendoorn, G. Tamminga, S. P. Hoogendoorn, and W. Daamen, "Longitudinal driving behavior under adverse weather conditions: Adaptation effects, model performance and freeway capacity in case of fog," in Proc. 13th Int. IEEE Conf. Intell. Transp. Syst., Sep. 2010, pp. 450–455.

9. S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model," Nature Commun., vol. 11, no. 1, p. 4850, Sep. 2020.

10. S. Liu, J. Wang, and T. Fu, "Effects of lane width, lane position and edge shoulder width on driving behavior in underground urban expressways: A driving simulator study," Int. J. Environ. Res. Public Health, vol. 13, no. 10, p. 1010, Oct. 2016.

11. R. Nagel, "Unraveling in guessing games: An experimental study," Amer. Econ. Rev., vol. 85, no. 5, pp. 1313–1326, 1995.

12. Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes. (2026). International Journal of Research and Applied Innovations, 9(1), 13641-13649. https://doi.org/10.15662/IJRAI.2026.0901017

13. F. Rosey, I. Aillerie, S. Espié, and F. Vienne, "Driver behaviour in fog is not only a question of degraded visibility—A simulator study," Saf. Sci., vol. 95, pp. 50–61, Jun. 2017.

14. D. D. Salvucci, "Modeling driver behavior in a cognitive architecture," Hum. Factors, vol. 48, no. 2, pp. 362–380, Jun. 2006.

15. K. Sama, Y. Morales, H. Liu, N. Akai, A. Carballo, E. Takeuchi, and K. Takeda, "Extracting human-like driving behaviors from expert driver data using deep learning," IEEE Trans. Veh. Technol., vol. 69, no. 9, pp. 9315–9329, Sep. 2020.

16. Y. Shao, J. Xu, B. Li, and K. Yang, "Modeling the speed choice behaviors of drivers on mountainous roads with complicated shapes," Adv. Mech. Eng., vol. 7, no. 2, Jan. 2015, Art. no. 862610.

17. A. Taragin and L. Leisch, "Driver performance on horizontal curves," in Proc. Highway Res. Board, vol. 33, 1954, pp. 446–466.

18. R. Fuller, "Towards a general theory of driver behaviour," Accident Anal. Prevention, vol. 37, no. 3, pp. 461–472, May 2005.

19. J. A. Villacorta-Atienza, "Static internal representation of dynamic situations reveals time compaction in human cognition," J. Adv. Res., vol. 28, pp. 111–125, Feb. 2021.

20. Y. Wiseman and I. Grinberg, "Circumspectly crash of autonomous vehicles," in Proc. IEEE Int. Conf. Electro Inf. Technol. (EIT), May 2016, pp. 387–392.

21. C. Wu and Y. Liu, "Queuing network modeling of the psychological refractory period (PRP)," Psychol. Rev., vol. 115, no. 4, pp. 913–954, 2008.

22. D. Xu, "Learning from naturalistic driving data for human-like autonomous highway driving," IEEE Trans. Intell. Transp. Syst., vol. 22, no. 12, pp. 7341–7354, Dec. 2021.

23. S. Xie, J. Zheng, and J. Wang, "Cognition-inspired behavioural feature identification and motion planning ways for human-like automated driving vehicles," IET Intell. Transp. Syst., vol. 17, no. 4, pp. 754–766, Apr. 2023.