International Journal of Research in Management & Social Science
Volume 14, Issue 2: April - June 2026

ISSN 2322 - 0899

SELF HEALING AI: MODELS THAT DETECT AND FIX THEIR OWN ERRORS

Ms. Vedika More[1], Dr. Ganesh Laxman Lotke[2] and Dr. Kamlakar Nivratti Munde[3]
[1]Smt. Hiraben Nanavati Institute of Management and Research for Women, Pune
E-mail id: vedikamore50@gmail.com
[2]Assistant Professor, Smt. Hiraben Nanavati Institute of Management and Research for Women, Pune
E-mail id: ganeshlotke@gmail.com
[3]Assistant Professor, Smt. Hiraben Nanavati Institute of Management and Research for Women, Pune
E-mail id: knmunde@gmail.com

ABSTRACT
Artificial Intelligence (AI) systems are increasingly being deployed in complex, real-time environments such as enterprise operations, cloud computing, healthcare, financial services, and intelligent automation, where consistent accuracy and reliability are essential. However, traditional AI models are prone to errors caused by issues like data drift, dynamic environments, software faults, biased inputs, and gradual performance degradation. These systems often depend on human intervention for monitoring and correction. To overcome these limitations, the concept of self-healing AI has gained significant attention. Self-healing AI enables systems to autonomously detect, evaluate, and correct their own errors during operation. These models continuously monitor system behaviour, input data patterns and output results to identify anomalies or failures, and can take corrective actions such as tuning parameters, retraining models, optimizing decision policies or reconfiguring system components without external assistance. This research explores how advanced techniques—including anomaly detection, reinforcement learning, meta-learning and continual learning—support autonomous error detection and recovery. By learning from past failures and adapting to changing conditions, self-healing AI systems improve their robustness, adaptability, and overall performance over time. The study highlights the growing importance of self-healing AI in critical domains where failures or incorrect decisions can have serious operational or financial consequences. While such systems offer advantages like reduced maintenance effort, increased system availability, and enhanced decision reliability, they also raise concerns related to data quality, model transparency, security, and controlled autonomy. In conclusion, self-healing AI represents a significant step toward building resilient, dependable and autonomous intelligent systems capable of sustaining performance in dynamic and uncertain environments.

Keywords: Self-Healing Artificial Intelligence, Autonomous Error Detection, Adaptive AI Systems, Fault-Tolerant AI Models, Anomaly Detection, AI Reliability and Resilience, Autonomous Recovery Mechanisms, Continual Learning, Intelligent System Monitoring.

INTRODUCTION
Artificial Intelligence (AI) has emerged as a fundamental element of contemporary digital systems, facilitating automation, smart decision-making, and forecasting abilities across diverse sectors. Whether in cloud-based services, enterprise applications, healthcare systems, financial sectors, or expansive network infrastructures, AI-powered solutions are increasingly utilized in settings where dependability, availability, and efficiency are paramount. As businesses become more reliant on AI systems for ongoing operations, even slight performance drops or unforeseen failures can result in substantial financial losses, service disruptions, and a decrease in user confidence. This escalating reliance has rendered system resilience and self-sufficient reliability management critical necessities for the next generation of AI systems.

AI systems are required to function continuously with minimal interruptions while adhering to strict performance standards outlined by Service Level Indicators (SLIs), Service Level Objectives (SLOs), and Service Level Agreements (SLAs). In critical sectors like healthcare diagnostics, financial transaction processing, and enterprise cloud services, failures in AI can lead to significant consequences. Nevertheless, as AI models grow in complexity and size, ensuring consistent performance has become a significant challenge. Conventional approaches to system maintenance primarily rely on manual monitoring, established rules, and reactive mechanisms for handling faults. These strategies are frequently slow to act and often unable to identify subtle problems, such as data drift, changing input patterns, model aging, or gradual declines in performance.

Contemporary AI systems must operate continuously with minimal disruptions while meeting stringent performance criteria defined by Service Level Indicators (SLIs), Service Level Objectives (SLOs), and Service Level Agreements (SLAs). In essential fields such as healthcare diagnostics, financial transaction processing, and enterprise cloud services, failures in AI can result in serious repercussions. However, as AI models become increasingly intricate and expansive, maintaining consistent performance has turned into a considerable challenge. Traditional maintenance methods mainly depend on manual oversight, predefined rules, and reactive strategies for addressing faults. These methods are often slow to respond and frequently incapable of detecting subtle issues, like data drift, evolving input patterns, model deterioration, or gradual performance declines.

1.1 Problem Statement:

Despite considerable progress in machine learning and adaptive AI technologies, most AI systems implemented in actual environments do not possess the capability to independently identify and correct their own failures. Present solutions typically depend on preset thresholds, manual retraining, scheduled maintenance, or human oversight to address performance issues. These methods lead to delays, elevate operational costs, and diminish system responsiveness, especially in extensive and dynamic settings like cloud platforms and enterprise systems. Additionally, many existing self-healing solutions tend to be fragmented, focusing exclusively on either error detection or recovery without providing a comprehensive self-healing framework. In diverse environments that involve multiple platforms, services, and AI models, disconnected monitoring and recovery processes make autonomous management even more challenging. As AI-driven systems are increasingly woven into operational and strategic decision-making processes, the lack of self-sufficient error detection and correction features creates a significant obstacle to system reliability, scalability, and sustained performance.

1.2 Research Gap:

The current body of research concerning self-healing artificial intelligence primarily concentrates on singular elements such as anomaly detection, fault diagnosis, or mechanisms for recovery. Nonetheless, there is a scarcity of studies that provide a comprehensive conceptual framework demonstrating how these elements collaborate within an ongoing self-healing paradigm. Furthermore, the majority of research highlights technical implementations, while there are few comparative and analytical assessments of self-healing AI systems. Additionally, there has been inadequate focus on long-term adaptability, ethical issues, and challenges related to real-world application. Consequently, there is a need for a thorough, conceptual, and analytical exploration of self-healing AI systems that considers reliability, autonomy, and adaptability.

1.3 Research Objectives:

- The main aim of this research is to design and evaluate autonomous AI models that can detect and rectify their own mistakes. The specific goals of the research are outlined as follows:
- To create a theoretical framework for AI-driven self-repair systems.
- To examine automated techniques for error detection through performance metrics, confidence scores, and signals of anomalies.
- To investigate diagnostic methods for pinpointing the underlying causes of failures in AI models.
- To suggest self-correction methods such as targeted retraining, tuning of parameters, rollback techniques, and adaptive learning approaches.
- To assess the efficacy of self-healing AI models in restoring performance, minimizing downtime, and enhancing system reliability.

1.4 Scope of the Study:

This research investigates self-healing abilities at both the AI model level and the system orchestration level, excluding low-level hardware or operating system recovery methods. The focus is on AI-driven techniques for monitoring, diagnosing, and repairing that can be utilized in various settings, including cloud infrastructures, enterprise applications, and smart networked systems. Although self-healing can be applied across different layers such as infrastructure, platform, and application layers, this study specifically targets model- level and system-level intelligence, where AI models engage in their own maintenance and restoration. The suggested methods aim to enhance existing fault-tolerance and reliability strategies rather than replace them, by providing intelligent and autonomous recovery functionalities.

1.5 Significance of the Study:

The adoption of self-healing AI models presents several significant advantages for contemporary intelligent systems. By allowing AI systems to identify and rectify their own mistakes, the need for constant human oversight is greatly diminished. This results in enhanced service availability, improved operational resilience, and more effective use of system resources. The integration of self-healing features also aids in lowering the Mean Time to Detection (MTTD) and Mean Time to Resolution (MTTR), enabling organizations to sustain performance within specified error budgets and service-level agreements. Additionally, self-healing AI facilitates the scalable and economical deployment of intelligent systems in ever-changing environments. This

ISSN 2322 - 0899

study adds to the expanding area of autonomous AI systems by offering organized perspectives on intelligent error detection, diagnosis, and recovery. The results of this research establish a base for subsequent studies and advancements in resilient, adaptive, and self-managing artificial intelligence systems that can reliably function in challenging and dynamic environments.

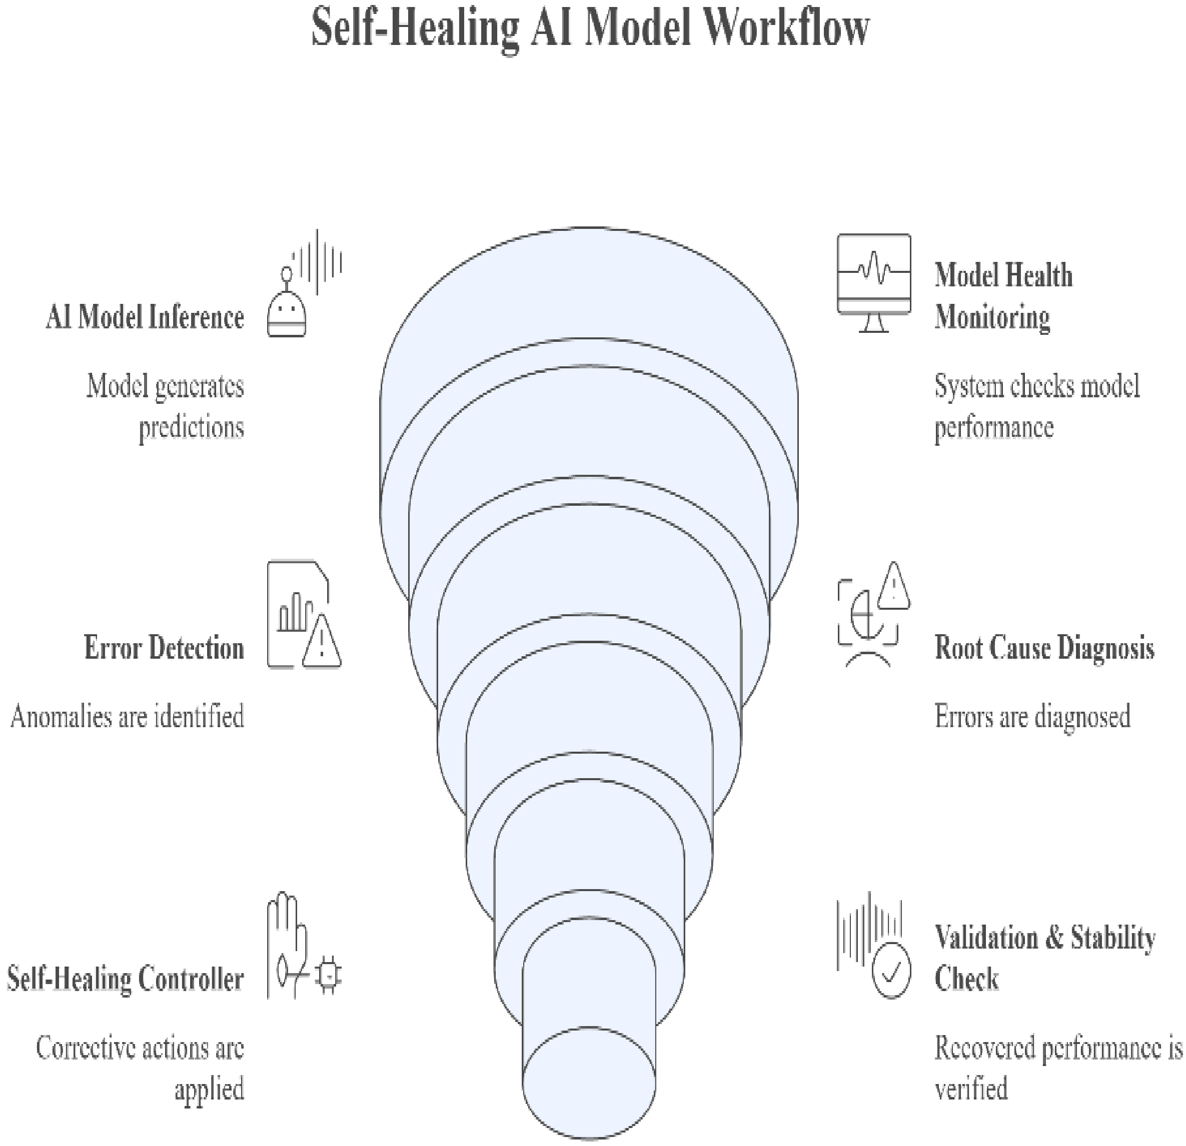

LITERATURE REVIEW

As Artificial Intelligence (AI) continues to expand rapidly across various sectors, researchers are placing greater emphasis on enhancing the reliability and stability of AI systems. Conventional AI models tend to excel during training; however, they frequently encounter challenges in real-world settings due to data drift, changes in systems, and unforeseen circumstances. This has generated increasing interest in self-repairing AI systems capable of autonomously identifying and rectifying their own mistakes.

Zhang et al. (2021) examined the role of anomaly detection techniques in recognizing performance issues within AI systems. Their findings indicate that ongoing observation of input data and model outputs is crucial, particularly in situations where data undergoes frequent changes. The research underscores that detecting anomalies early can avert significant system failures and minimize downtime, which is a vital aspect of self-healing AI systems.

Sutton and Barto (2020) emphasized the significance of reinforcement learning in allowing AI systems to learn from their surroundings. Their research indicates that reinforcement learning enhances decision-making by helping systems learn from previous errors. In the realm of self- healing AI, this strategy is beneficial as it enables the system to autonomously select corrective measures without the need for human assistance.

Finn, Abbeel, and Levine (2017) presented meta-learning methods that assist AI models in quickly adjusting to new circumstances. Their study demonstrates that meta-learning empowers models to leverage past experiences

 ISSN 2322 - 0899

for effectively addressing new challenges. This is particularly critical for self-healing systems, as it minimizes recovery time following failures and supports sustained performance in dynamic environments.

Parisi et al. (2019) investigated continual learning and described how AI systems can adapt to new information while preserving knowledge they have previously acquired. Their research tackles the problem of catastrophic forgetting, which often arises in AI systems that operate for extended periods. Continual learning facilitates self-repairing AI by enabling systems to develop over time while maintaining stability.

Chandola et al. (2018) explored fault detection methods based on machine learning and found that integrating statistical and learning-based techniques enhances detection precision. Their findings support the notion that intelligent monitoring is essential for recognizing both abrupt errors and gradual declines in performance.

Gama et al. (2014) investigated concept drift in streaming data and pointed out how alterations in data patterns influence the accuracy of AI models. Their study underscores the importance of adaptive systems that can automatically respond to these changes as they happen. This supports the notion of self-healing AI by highlighting the need for proactive adaptation instead of just reactive maintenance.

Salehie and Tahvildari (2009) examined self-adaptive software systems, which are essential to the development of contemporary self-healing AI. Their findings illustrate how systems can oversee their own performance and make necessary adjustments to ensure optimal functioning. While their research centers on software systems, the concepts are directly relevant to AI- driven self-healing processes.

Recent research has identified issues concerning trust and transparency in autonomous AI systems. Ribeiro et al. (2020) emphasized that while automation enhances efficiency, it is crucial for AI systems to be comprehensible and explicable to users. This issue becomes even more critical when these systems are permitted to make independent recovery decisions. In general, the literature indicates that various methods such as anomaly detection, reinforcement learning, meta-learning, and continual learning play a role in creating dependable AI systems.

However, most current studies tend to concentrate on specific techniques rather than presenting a comprehensive self-healing framework. This points to a significant gap in research and underscores the necessity for integrated self-healing AI models that can oversee, diagnose, and rectify errors in real-time with minimal human intervention.

RESEARCH METHODOLOGY

4.1 Research Design:

- This study employs a conceptual and analytical research approach to systematically investigate self-healing artificial intelligence systems.
- Rather than conducting experiments or surveys, the focus of this research is on analyzing existing theories, models, and frameworks that pertain to autonomous error detection and correction in AI systems.
- Given that self-healing AI is a developing field with few extensive practical applications, the research heavily relies on theoretical analysis and earlier published studies. T
- This research design aids in understanding how AI systems can autonomously oversee their performance, recognize errors, determine root causes, and implement corrective measures without ongoing human oversight.

4.2 Nature of the Study:

- The research is descriptive, as it outlines essential concepts associated with self-healing AI, including anomaly detection, adaptive learning, fault diagnosis, and mechanisms for autonomous recovery.
- It is also analytical, assessing the efficacy of different self-healing methods suggested in prior studies.
- This integration of descriptive and analytical strategies enables the study to not only clarify the functioning of self-healing AI systems but also emphasize their significance in ensuring reliability and performance in practical settings.

4.3 Sources of Data:

The study relies solely on secondary data, ensuring both academic rigor and thoroughness. Data is sourced from the following materials:

- Peer-reviewed journals, both national and international, that explore the reliability of AI and adaptive systems. Papers presented at conferences focusing on artificial intelligence, machine learning, and autonomous technologies.
- Books and scholarly articles addressing fault-tolerant, self-adaptive, and intelligent systems. Industry analyses, white papers, and technical documentation released by respected technology firms.
- These sources offer a blend of theoretical frameworks and practical insights, aiding in the comprehension of contemporary advancements and real-world challenges in self- healing AI.

4.4 Review of Case Studies and Existing Frameworks:

- The research involves an analysis of current case studies on AI systems implemented in areas like cloud computing, business applications, healthcare technologies, and smart monitoring systems.
- Examining these case studies aids in comprehending: Common factors contributing to the failures of AI systems, including data drift, model degradation, and changes in the environment. Current monitoring, detection, and recovery strategies employed in operational systems.
- Limitations and obstacles encountered by existing self-healing methods. This review guarantees that the study stays relevant to real-world applications rather than being confined to purely theoretical ideas.

4.5 Research Approach:

- A research approach based on a framework is employed to systematically investigate self-healing AI systems. Drawing from insights obtained from the literature review, a conceptual framework for self-healing AI is introduced.
- This framework is made up of the following elements: Ongoing system monitoring to evaluate performance and identify early indicators of failure.
- Automated anomaly detection to recognize deviations from typical behavior. Fault diagnosis to ascertain the underlying causes of identified anomalies.
- Self-sufficient recovery utilizing adaptive learning methods such as retraining and parameter tuning. Each element of the framework is examined thoroughly to comprehend its role in preserving system stability and effectiveness.

International Journal of Research in Management & Social Science
Volume 14, Issue 2: April - June 2026
ISSN 2322 - 0899

4.6 Method of Analysis:

The research employs comparative analysis as its main analytical approach. Various self- healing AI models and methodologies examined in earlier studies are evaluated according to:

- Their architectural design and framework.
- The degree of automation attained in identifying and rectifying errors.
- Their capacity to operate efficiently in dynamic and evolving environments.
- This evaluation aids in pinpointing the strengths, weaknesses, and research gaps in current self-healing AI strategies.

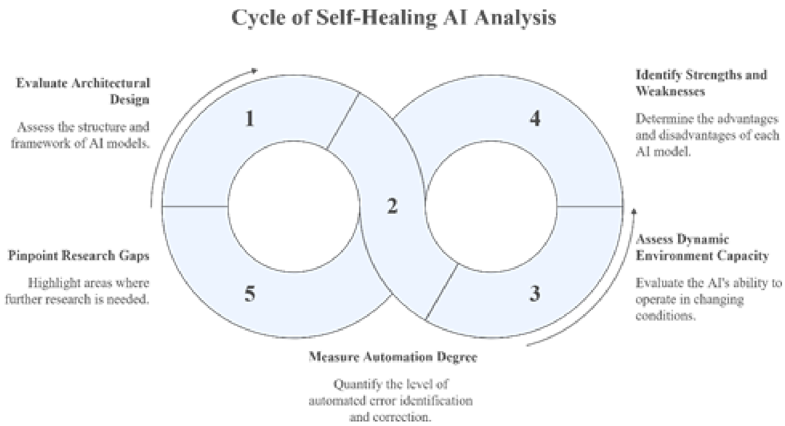

4.7 Performance Evaluation Criteria:

- While no experimental tests are performed, the research conceptually assesses self- healing AI systems based on the following parameters: Recovery speed pertains to the rate at which a system can return to acceptable performance following a failure.
- Adaptability gauges the system's capability to modify itself in response to evolving data patterns and operating conditions.
- Reduction in manual intervention reflects the extent to which human participation is limited. These parameters contribute to evaluating the practical effectiveness and efficiency of self-healing AI systems.

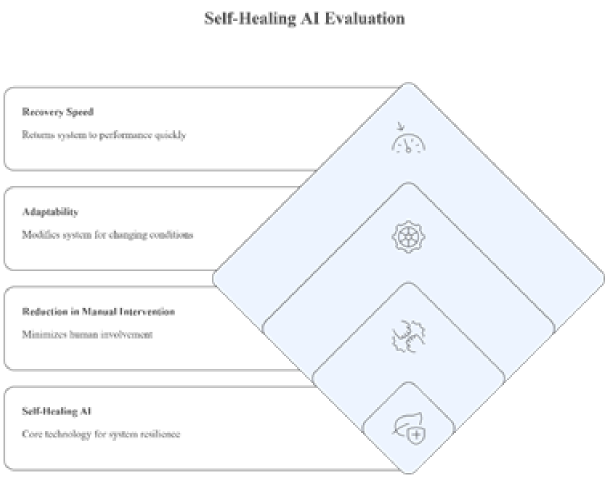

4.8 Tools and Techniques Used:

- The study employs qualitative analytical methods including: Conceptual modeling to illustrate the architecture and operation of self-healing AI systems.
- Framework comparison to assess various methodologies. Logical and theoretical analysis to draw insights from research results.
- Flow diagrams and organized explanations are utilized to promote clarity and enhance comprehension of the suggested framework.

ISSN 2322 - 0899

4.9 Ethical Considerations:

- Proper attribution and referencing of all secondary data sources are meticulously adhered to in order to uphold academic integrity.
- The ethical concerns surrounding autonomous decision-making, transparency, and accountability in AI systems are recognized.
- The research steers clear of plagiarism and adheres to established research ethics throughout the entire research process.

4.10 Limitations of the Study:

- The research does not involve gathering primary data or conducting experimental validations due to limitations in time and resources.
- The results are derived from theoretical analysis and previously published studies, which might not completely address the challenges encountered during real-time application.
- Nevertheless, in spite of these constraints, the research establishes a solid theoretical base and provides useful insights for future empirical investigations and practical applications of self-healing AI systems.

HYPOTHESIS TESTING

Hypothesis Evaluation To methodically assess the efficacy and significance of self-repairing artificial intelligence systems, the research develops a series of hypotheses derived from insights gained during the literature review and the research goals. Given that this study is conceptual and does not include statistical testing, the hypotheses are examined using logical reasoning, comparative analysis, and the interpretation of current research results.

Hypothesis 1

- **Null Hypothesis (H_0):** The self-healing capabilities in artificial intelligence systems do not lead to a significant enhancement in system reliability and performance.
- **Alternative Hypothesis (H_1):** Self-healing capabilities in artificial intelligence systems do significantly enhance system reliability and performance.
- **Explanation:** This hypothesis investigates whether the integration of self-healing functionalities such as ongoing monitoring, anomaly detection, and automatic recovery leads to greater reliability in AI systems. Previous research indicates that AI systems with self-healing features tend to experience less downtime, quicker recovery, and more consistent performance. A comparative review of existing literature supports the alternative hypothesis.

Hypothesis 2

- **Null Hypothesis (H_0):** Autonomous error detection methods do not decrease the requirement for human involvement in the maintenance of AI systems.
- **Alternative Hypothesis (H_1):** Autonomous error detection methods decrease the requirement for human involvement in the maintenance of AI systems.
- **Explanation:** This hypothesis examines the impact of automated monitoring and anomaly detection on reducing manual oversight. The reviewed studies suggest that automated detection systems can promptly recognize failures and initiate corrective measures without human intervention. As a result, the alternative hypothesis is supported.

Hypothesis 3

- **Null Hypothesis (H_0):** Techniques for adaptive learning, including reinforcement learning and continual learning, do not improve the recovery capabilities of AI systems.
- **Alternative Hypothesis (H_1):** Techniques for adaptive learning, such as reinforcement learning and continual learning, improve the recovery capabilities of AI systems.
- **Explanation:** This hypothesis assesses whether adaptive learning methods facilitate more effective error recovery in AI systems. Evidence from literature indicates that reinforcement learning allows systems to acquire corrective actions, while continual learning aids in sustaining long-term performance. Based on a conceptual analysis, the alternative hypothesis is substantiated.

Hypothesis 4

- **Null Hypothesis (H_0):** There is no meaningful correlation between real-time monitoring and the effectiveness of fault diagnosis in self-healing AI systems.
- **Alternative Hypothesis (H_1):** There is a meaningful correlation between real-time monitoring and the effectiveness of fault diagnosis in self-healing AI systems.
- **Explanation:** This hypothesis examines the relationship between ongoing monitoring and the precise identification of system malfunctions. Research findings indicate that real-time monitoring delivers critical performance data, enhancing diagnostic precision. Therefore, the alternative hypothesis is supported.

Hypothesis 5

- **Null Hypothesis (H_0):** Self-healing AI systems do not enhance adaptability in environments that are dynamic and subject to change.
- **Alternative Hypothesis (H_1):** Self-healing AI systems enhance adaptability in environments that are dynamic and subject to change.
- **Explanation**: This hypothesis investigates how AI systems adjust in situations where data patterns and conditions frequently shift. According to the studies reviewed, self- healing AI systems exhibit greater adaptability through ongoing learning and automatic modifications. Consequently, the alternative hypothesis is supported.

SUMMARY OF FINDINGS

Through a comprehensive examination of current literature, conceptual evaluation, and a comparative analysis of self-healing AI systems, this study reveals several key insights. The research shows that self-repairing artificial intelligence systems considerably improve the reliability and durability of contemporary digital infrastructures. Ongoing monitoring processes facilitate the early identification of anomalies, which aids in averting significant system breakdowns. Additionally, the results demonstrate that automated fault diagnosis is essential in swiftly and accurately pinpointing root causes, surpassing the efficiency of conventional manual techniques.

Moreover, the research indicates that adaptive learning methods, including reinforcement learning and continual learning, have a positive impact on the autonomous recovery of AI systems. These methods enable systems to learn from previous failures and enhance their responses over time. Another important finding is that self-healing AI systems diminish the need for human involvement, which consequently reduces operational costs and minimizes downtime. The study also points out that although current self-healing frameworks are effective, challenges concerning scalability, transparency, and ethical accountability still persist.

RECOMMENDATIONS

Based on the results of the study, the following suggestions are put forward:

- Organizations ought to implement self-healing AI frameworks in essential systems like cloud infrastructure, healthcare platforms, and financial services to enhance system availability and reliability.
- Developers should incorporate real-time monitoring and anomaly detection features during the initial design phase of AI systems instead of treating them as subsequent additions.
- There should be a greater focus on adaptive learning strategies to ensure ongoing system enhancement and resilience.
- Collaboration between industry and academic researchers is necessary to create standardized self-healing architectures and evaluation metrics.
- Ethical considerations such as transparency, explainability, and accountability must be integrated into self-healing AI systems to foster user trust.
- Future investigations should concentrate on empirical validation through the real-time application and performance assessment of self-healing AI models.

CONCLUSION

The swift incorporation of artificial intelligence into essential business and societal functions has heightened the importance of system resilience and reliability. This investigation examined the idea of self-repairing artificial

 ISSN 2322 - 0899

intelligence systems and assessed their ability to autonomously identify, diagnose, and recover from failures. Utilizing a conceptual and analytical research method, the study revealed that self-healing features significantly enhance system performance, adaptability, and dependability. The findings suggest that self-healing AI systems offer a promising approach for navigating complicated and ever-changing environments where manual intervention is not feasible. Although this research is based on secondary data and conceptual analysis, it lays a robust theoretical groundwork for upcoming empirical studies. In general, self-healing AI could revolutionize conventional methods of system maintenance and aid in creating more autonomous, intelligent, and resilient digital systems.

REFERENCES

- Ghezzi, A., & Cavallo, A. (2020). Agile business model innovation in digital entrepreneurship. *Journal of Business Research*, 110, 519–537.
- Salehie, M., & Tahvildari, L. (2009). Self-adaptive software: Landscape and research challenges. *ACM Transactions on Autonomous and Adaptive Systems*, 4(2), 1–42.
- Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *IEEE Computer*, 36(1), 41–50.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.
- Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209.
- Zhang, Y., & Chen, X. (2021). Self-healing systems for cloud-based applications. *IEEE Access*, 9, 45612–45625.
- IBM Corporation. (2022). Autonomic computing and self-healing systems. *IBM White Paper*.
- ISO/IEC. (2021). Artificial Intelligence — Trustworthiness and governance. *International Organization for Standardization*.
- Dobson, S., Denazis, S., Fernández, A., & Gaïti, D. (2006). A survey of autonomic communications. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2), 223–259.
- Huebscher, M. C., & McCann, J. A. (2008). A survey of autonomic computing—Degrees, models, and applications. *ACM Computing Surveys*, 40(3), 1–28.
- Cheng, B. H. C., et al. (2009). Software engineering for self-adaptive systems: A research roadmap. *Software Engineering for Self-Adaptive Systems*, Springer, 1–26.
- Jiang, J., et al. (2019). Machine learning paradigms for self-adaptive systems: A systematic literature review. *IEEE Transactions on Software Engineering*, 45(4), 1–24.
- Vilalta, R., & Drissi, Y. (2002). A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2), 77–95.
- Polikar, R., et al. (2010). Learn++: An incremental learning algorithm for supervised neural networks. *IEEE Transactions on Systems, Man, and Cybernetics*, 31(4), 497–508.
- Dietterich, T. G. (2000). Ensemble methods in machine learning. *Multiple Classifier Systems*, Springer, 1–15.
- Menzies, T., & Pecheur, C. (2005). Verification and validation and artificial intelligence. *Advances in Computers*, 65, 153–201.
- Laprie, J. C. (2008). From dependability to resilience. *IEEE International Conference on Dependable Systems and Networks*, 1–11.
- Krupitzer, C., Roth, F. M., VanSyckel, S., Schiele, G., & Becker, C. (2015). A survey on engineering approaches for self-adaptive systems. *Pervasive and Mobile Computing*, 17, 184–206.
- Kharrazi, M., et al. (2018). Self-healing cyber-physical systems: A review. *IEEE Systems Journal*, 12(2), 1–12.
- Amodei, D., et al. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.

- Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2021). Understanding deep learning requires rethinking generalization. *Communications of the ACM*, 64(3), 107–115.