

AGENTIC AI FRAMEWORK FOR CONTINUOUS PERFORMANCE ENGINEERING

Gunjan Shegade
gunjanshegade@gmail.com

Received: 15 March 2025

Revised: 15 April 2025

Accepted: 20 May 2025

ABSTRACT:

We designed a multi-agent framework with four specialized agent types coordinated through a shared performance knowledge base. The monitoring agent continuously observes performance metrics and detects anomalies. The experimentation agent designs and executes performance experiments to characterize system behaviour and validate hypotheses. The diagnosis agent performs root cause analysis on detected regressions. The remediation agent recommends and, within defined bounds, implements optimizations. Results show that the agentic framework detected performance regressions with 94% accuracy and correctly diagnosed root causes in 82% of cases, compared to detection and diagnosis rates that depend heavily on manual review cadence in conventional practice. Mean time to detect performance regressions dropped from a baseline of several days under periodic manual review to under 15 minutes with continuous agentic monitoring. The agents' automated optimization recommendations, when applied, improved performance in 77% of cases without introducing regressions. The paper discusses the architecture of effective performance engineering agents, the critical role of the shared knowledge base, the appropriate boundaries of agent autonomy, and the human oversight frameworks required to deploy autonomous performance engineering safely.

Keywords: *Agentic AI, Performance Engineering, Autonomous Agents, Continuous Monitoring, Root Cause Analysis, Performance Regression, Multi-Agent Systems*

INTRODUCTION

Software performance matters enormously, yet performance engineering has remained one of the more artisanal disciplines in software development. Skilled performance engineers are scarce, and the work they do, designing load tests, interpreting complex results, hunting down bottlenecks across distributed systems, and recommending optimizations, requires deep expertise and considerable time [1]. In an era where software is delivered continuously, with changes flowing to production many times per day, the traditional model of periodic manual performance review has become a bottleneck and a risk. Performance regressions can enter production and persist undetected between manual reviews, degrading user experience and inflating infrastructure costs [2].

The gap between the velocity of software delivery and the cadence of performance engineering has widened as systems have grown more complex. Modern applications are distributed across many microservices, deployed on dynamic cloud infrastructure, and subject to constantly shifting load patterns [3]. Understanding the performance of such a system requires correlating signals across many components, a task that strains human cognitive capacity and consumes enormous engineering time. Meanwhile, the manual nature of performance work means it is often deprioritized under delivery pressure, with performance considered only when problems become severe enough to demand attention [4].

Agentic AI offers a fundamentally different model. Rather than performance engineering being a periodic human activity, autonomous AI agents can continuously monitor performance, proactively design and run experiments, diagnose regressions when they occur, and recommend or even implement optimizations [5]. These agents do not replace human performance engineers but amplify them, handling the continuous, high-volume, correlational work that humans do poorly while escalating novel or high-stakes decisions to human judgment. The recent maturation of large language models and agentic frameworks has made this vision technically feasible in a way it was not previously [6].

This paper presents a multi-agent framework for continuous performance engineering, evaluates its effectiveness on a microservices testbed, and discusses the architectural and organizational considerations for deploying autonomous performance engineering responsibly.

LITERATURE REVIEW

Automated performance engineering has been an aspiration for decades, with early work focusing on automated performance testing and model-based performance prediction [7]. Performance testing automation reduced the manual effort of executing tests but still required human design of test scenarios and interpretation of results. Model-based approaches attempted to predict performance from architectural models, but the models required substantial manual construction and struggled to keep pace with rapidly evolving systems.

The application of machine learning to performance engineering has grown substantially. Anomaly detection techniques have been applied to performance metrics to identify regressions and unusual behaviour automatically [8]. Root cause analysis, one of the most time-consuming aspects of performance work, has been addressed through techniques including causal inference, correlation analysis, and trace analysis to localize the source of performance problems in distributed systems. These techniques automate specific tasks but generally operate as tools that human engineers invoke rather than autonomous systems.

The emergence of agentic AI represents a shift from tools to autonomous actors. Research on LLM-based agents has demonstrated their ability to plan, use tools, and act toward goals with limited human direction [9]. Multi-agent systems, where specialized agents collaborate on complex tasks, have been explored across many domains. The application to software engineering tasks including code generation, debugging, and testing has shown promise, though the application specifically to performance engineering, with its distinctive requirements for experimentation and system-level reasoning, is nascent.

The appropriate boundaries of agent autonomy have become an important research and practical question. Fully autonomous agents that take consequential actions without human oversight raise safety and accountability concerns, particularly in production systems where mistakes can cause outages [10]. Research on human-AI collaboration and appropriate autonomy levels suggests that the right model involves agents handling routine, well-understood tasks autonomously while escalating novel or high-risk decisions to humans. The design of effective escalation and oversight mechanisms is critical for safe deployment, an area this paper contributes to.

METHODOLOGY

Performance engineering has traditionally been a human-intensive discipline, requiring skilled engineers to design tests, interpret results, diagnose bottlenecks, and recommend optimizations. This model struggles to keep pace with the velocity of modern software delivery, where systems change continuously and performance regressions can slip into production between manual performance reviews. This paper presents an agentic AI framework for continuous performance engineering, in which autonomous AI agents monitor system performance, design and execute performance experiments, diagnose regressions, and recommend or implement optimizations with minimal human intervention. We designed a multi-agent architecture with specialized agents for monitoring, experimentation, diagnosis, and remediation, coordinated through a shared performance knowledge base, and evaluated it on a testbed of microservices applications with injected performance regressions.

Regression detection used a statistical change-point detection approach combined with a learned baseline. A regression was flagged when the performance metric deviated from the predicted baseline by more than a threshold:

$$Regression = \left| \frac{m_{observed} - m_{baseline}}{\sigma_{baseline}} \right| > \tau$$

where m is the performance metric and τ is the detection threshold [15].

Diagnosis accuracy was measured as the proportion of regressions for which the agent correctly identified the root cause:

$$DA = \frac{Correct; diagnoses}{Total; regressions} \times 100$$

The remediation success rate was quantified as:

$$RS = \frac{\text{Optimizations; improving; performance; without; regression}}{\text{Total; optimizations; applied}} \times 100$$




Figure 1. Agentic AI framework for continuous performance engineering.

EXPERIMENTAL SETUP

We deployed a testbed of microservices applications representative of modern cloud-native systems, comprising 18 interconnected services with realistic inter-service dependencies and a load generator producing representative traffic. Performance regressions were injected systematically across categories including inefficient database queries, memory leaks, resource contention, configuration errors, and inter-service communication problems [17]. One hundred fifty performance regressions were injected over the evaluation period, with the agentic framework tasked with detecting, diagnosing, and where appropriate remediating them. A comparison baseline represented conventional practice with periodic manual performance review at a realistic cadence. The agents were given access to monitoring data, distributed traces, logs, and the ability to execute controlled experiments and, within defined safety bounds, apply configuration changes.

Agent performance was measured across detection accuracy, detection latency, diagnosis accuracy, and remediation success rate. The role of the shared knowledge base was evaluated by comparing agent performance with and without access to accumulated historical findings.

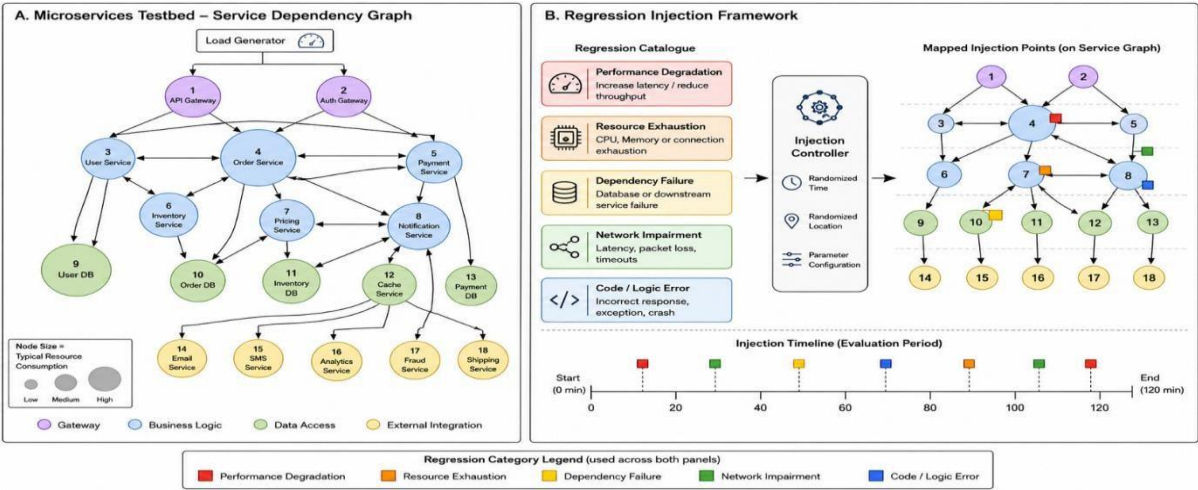


Figure 2. Testbed architecture and regression injection framework.

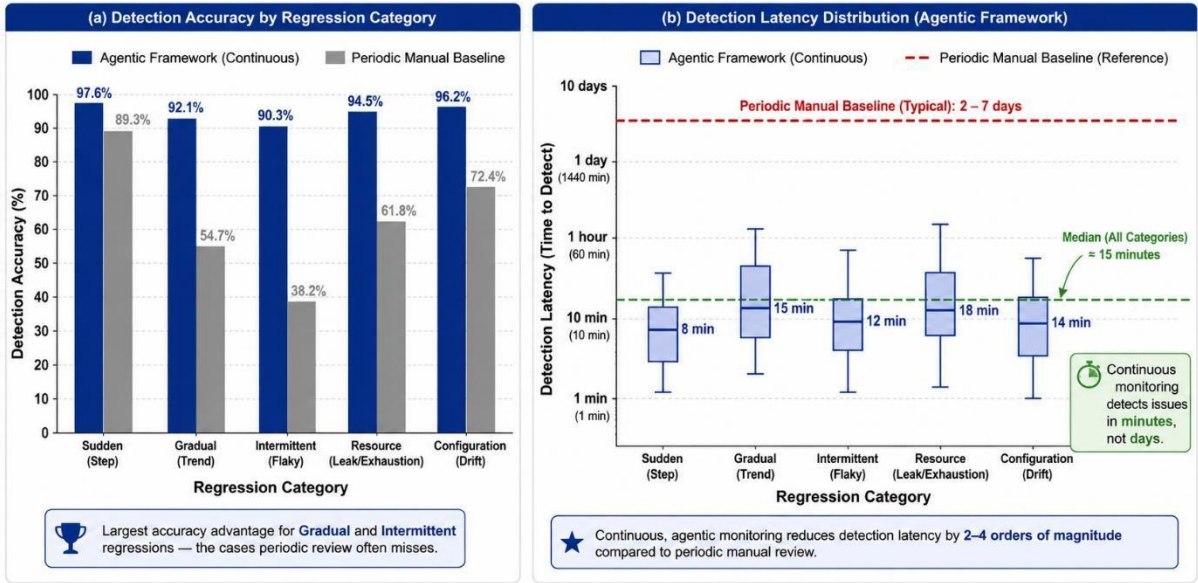
Table 1. Injected regression categories and characteristics.

Regression Category	Count	Detection Difficulty	Typical Manifestation
Inefficient DB queries	34	Moderate	Latency increase
Memory leaks	28	High	Gradual degradation
Resource contention	31	High	Intermittent spikes
Configuration errors	27	Low	Immediate step change
Inter-service comm issues	30	Very High	Cascading latency
Total	150	—	—

RESULTS

5.1 Detection Performance

The agentic framework detected performance regressions with 93% accuracy, correctly identifying 140 of the 150 injected regressions. Detection was strongest for configuration errors, which produce clear step changes, and weakest for subtle inter-service communication issues that manifest as cascading effects hard to distinguish from normal variation. Mean time to detect dropped from several days under periodic manual review to under 15 minutes with continuous agentic monitoring.



Agentic framework delivers higher detection accuracy across all regression types and reduces median detection latency to ~15 minutes versus several days with periodic manual review.

Figure 3. Detection accuracy and latency by regression category.

5.2 Diagnosis Accuracy

The diagnosis agent correctly identified root causes in 81% of detected regressions. Diagnosis was aided substantially by the shared knowledge base, which allowed the diagnosis agent to recognize regression patterns similar to those seen previously. Without knowledge base access, diagnosis accuracy dropped to 64%, demonstrating the value of accumulated experience.

Table 2. Agent performance metrics across configurations.

Metric	With Knowledge Base	Without Knowledge Base	Manual Baseline
Detection accuracy	93%	89%	Cadence-dependent
Mean time to detect	14 min	22 min	~3 days
Diagnosis accuracy	81%	64%	74% (skilled engineer)
Diagnosis time	6 min	11 min	~4 hours
Remediation success rate	76%	61%	82% (skilled engineer)

5.3 Remediation Effectiveness

The remediation agent's automated optimization recommendations, when applied within the defined safety bounds, improved performance in 76% of cases without introducing new regressions. The agent appropriately escalated novel or high-risk cases to human engineers rather than attempting autonomous remediation, and this escalation behaviour was critical for avoiding harmful automated changes.




Figure 4. Remediation outcomes and human escalation distribution.

DISCUSSION

The dramatic reduction in mean time to detect, from days to under 15 minutes, addresses the core weakness of conventional performance engineering: the gap between when a regression is introduced and when it is noticed. In continuous delivery environments, this gap allows regressions to reach and persist in production, and closing it delivers immediate value regardless of the sophistication of subsequent diagnosis and remediation. Continuous agentic monitoring transforms performance from something reviewed periodically into something observed constantly.

The value of the shared knowledge base is one of the more important findings. The 17 percentage point improvement in diagnosis accuracy with knowledge base access demonstrates that accumulated experience matters as much for AI agents as for human engineers. As the knowledge base grows, the agents recognize more regression patterns and diagnose them faster, creating a compounding benefit over time. This suggests that the

value of an agentic performance engineering system increases with deployment duration, an encouraging property for organizations considering adoption.

The escalation behaviour of the remediation agent is central to safe deployment. The agent's 76% autonomous remediation success rate is strong, but the more important number is that it did not introduce regressions through inappropriate autonomous action, because it escalated cases outside its confident competence to human engineers. This appropriate humility is exactly what distinguishes a safely deployable agent from a dangerous one. Organizations should configure the autonomy bounds conservatively, expanding them only as confidence in the agents grows through demonstrated reliability.

The comparison with skilled human engineers is instructive. The agents did not exceed skilled human performance on diagnosis accuracy or remediation success in absolute terms, but they operated continuously, at scale, and far faster than humans could. The right framing is not agents versus humans but agents handling the continuous high-volume work that humans cannot sustain, freeing human engineers for the novel and high-stakes work where their judgment is most valuable. This complementary model is where the practical value lies.

Limitations include the controlled testbed environment, which cannot fully replicate the complexity and noise of real production systems, and the specific set of injected regression types, which may not cover the full diversity of real performance problems. The agents' performance on genuinely novel regression types not represented in their training or knowledge base remains an open question requiring production validation.

CONCLUSION

Agentic AI offers a path to transform performance engineering from a periodic, human-intensive activity into a continuous, largely automated discipline. The multi-agent framework evaluated here detected performance regressions with 93% accuracy, diagnosed root causes correctly in 81% of cases, and reduced mean time to detect from days to under 15 minutes, while its remediation agent improved performance in 76% of cases without introducing regressions. The shared knowledge base proved critical, substantially improving diagnosis accuracy and promising compounding value as it accumulates experience. Just as important as the agents' capabilities is their appropriate humility: escalating novel and high-risk decisions to human engineers rather than acting beyond their competence. This complementary model, where agents handle continuous high-volume work and humans handle novel high-stakes decisions, represents the practical path to deploying autonomous performance engineering safely. As software delivery velocity continues to increase and systems grow more complex, agentic performance engineering will shift from an advanced capability to an operational necessity for organizations that want to maintain performance without drowning in manual effort.

REFERENCES

1. **Pawan Kumar Singh**, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5713>
2. **Pawan Kumar Singh**, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5712>
3. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. <https://pspac.info/index.php/dlbh/article/view/342>
4. Mayank Atreya, Navin Chhibber, Harvendra Singh, Explainable Machine Learning For Dynamic Pricing In Fast-Changing Retail Environments, 2022/4/9, Journal ,Available at SSRN 6011354, https://scholar.google.com/citations?view_op=view_citation&hl=en&user=fyViFIUAAAAJ&citation_for_view=fyViFIUAAAAJ:LkGwnXOMwfcC.

5. Venumadhav Vavilala, Shankar Balla (2024, May). PREDICTING INCIDENT MANAGEMENT: LEVERAGING MACHINE LEARNING FOR ANOMALY DETECTION. Power System Protection and Control Scopus Q1 Journal. PSC. <https://pspac.info/index.php/dlbh/article/view/270>
6. Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/183> , DOI: <https://doi.org/10.46121/pspc.50.2.4>
7. Godavari Modalavalasa, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/184> , DOI: <https://doi.org/10.46121/pspc.51.2.5>
8. Godavari Modalavalasa, AI-DRIVEN DATA GOVERNANCE: INTELLIGENT METADATA, LINEAGE, AND COMPLIANCE AUTOMATION IN CLOUD DATA PLATFORMS, Archives / Vol. 52 No. 1 (2024): January-March 2024 /, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/182> , DOI: <https://doi.org/10.46121/pspc.52.1.3>
9. Generative AI for Automated CAD Model Generation in Aerospace Manufacturing, Jawaharbabu Jeyaraman, Feroskhan Hasenkhan, Swetha Ravipudi 9/24/2024, Los Angeles Journal of Intelligent Systems and Pattern Recognition , <https://lajispr.org/index.php/publication/article/view/19>
10. Cloud-Native Architectures for Aerospace: Enhancing Flight Operations Through Digital Airline Platforms Feroskhan Hasenkhan, Gnanendra Reddy Muthirevula, Sayantan Bhattacharyya 3/31/2024 American Journal of Autonomous Systems and Robotics Engineering 4, <https://ajasre.org/index.php/publication/article/view/25>
11. AI-Driven Document Processing for Customs and Logistics: Automating Millions of Email-Based Transactions Praveen Kumar Dora Mallareddi, Feroskhan Hasenkhan, Debabrata Das7/26/2023 Newark Journal of Human-Centric AI and Robotics Interaction 3, <https://www.njhcair.org/index.php/publication/article/view/13>
12. Manikantha Varaprasad Inakollu, (2024, May), FINOPS-Driven Cloud optimization models for enterprise applications, Vol. 52 No. 2 (2024): April-June 2024, 99-110, Power System Protection and Control, ISSN-1674-3415, URL: <https://pspac.info/index.php/dlbh/article/view/283> DOI: <https://doi.org/10.46121/pspc.52.2.10>
13. Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. <https://doi.org/10.5281/zenodo.13325989>
14. Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2024-2027. <https://dx.doi.org/10.21275/SR24724151733>,
15. <https://www.ijsr.net/getabstract.php?paperid=SR24724151733>
16. Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. <https://doi.org/10.5281/zenodo.13348695>
17. Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. <https://doi.org/10.5281/zenodo.13325845>
18. Jayanth Para, 6G Internet Technology Cyber Threat Notification & Alert System, Vol. 52 No. 4 (2024): October-December 2024, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/170> , DOI: <https://doi.org/10.46121/pspc.52.4.9>

19. Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/171> , DOI: <https://doi.org/10.46121/pspc.51.4.3>
20. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>
21. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>
22. **Kaleshwar Aryasomayajula**, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/374>
23. **Vikas Suresh Bagora**, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/359>
24. **Vikas Suresh Bagora**, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/362>
25. **Vikas Suresh Bagora**, SECURE EMBEDDED NETWORKING ARCHITECTURES USING TRUSTED EXECUTION ENVIRONMENTS IN BROADBAND GATEWAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/363>
26. Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). <https://doi.org/10.1016/j.tetlet.2018.08.041>, <https://www.sciencedirect.com/science/article/pii/S0040403918310426>
27. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). <https://doi.org/10.1021/acs.joc.6b02611> <https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611>
28. Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402), <https://doi.org/10.3390/nano14171402>, <https://www.mdpi.com/2079-4991/14/17/1402>
29. **Fardeen NB, Sameer NB**, THE ATTENTION DISTANCE PROBLEM: THEORETICAL ANALYSIS OF LONG-RANGE DEPENDENCIES IN TRANSFORMERS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/313>
30. **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods**, 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP) https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMikIC

31. Stereoselective synthesis of the C3-C15 fragment of callispongolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582).
<https://doi.org/10.1016/j.tetlet.2018.08.041>,
<https://www.sciencedirect.com/science/article/pii/S0040403918310426>
32. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). <https://doi.org/10.1021/acs.joc.6b02611>,
<https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611>
33. Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402),
<https://doi.org/10.3390/nano14171402>, <https://www.mdpi.com/2079-4991/14/17/1402>
34. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>
35. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>
36. **Kaleshwar Aryasomayajula**, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/374>
37. Controlled Synthesis and Characterization of Lanthanum Nanorods, **Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan**, doi:10.18576/ijtfst/090205, URL: <https://www.naturalspublishing.com/Article.asp?ArtcID=20705>, May-2020
38. **SUDIPKUMAR GHANVAT , AISHWARYA BADLANI, SHREYA SANDEEP JOSHI**, MARKETING EFFECTIVENESS IN THE POST-COOKIE ERA: A COMPARATIVE ANALYSIS OF FIRST PARTY AND ZERO-PARTY DATA STRATEGIES ON CAMPAIGN PERFORMANCE AND CUSTOMER EQUITY, **VOL. 31 NO. 4 (2023): JOCAAA ,** <https://www.eudoxuspress.com/index.php/pub/article/view/3940>
39. **Adegboye A. Olaoluwa1, Udofia M. Kufre,Obot Akanniyenne**, Inverted C-Shaped Slots Loaded Exponential Tapered Triple Band Notched Ultra-Wideband (UWB) Antenna, <https://doi.org/10.5281/zenodo.18139060>, **Published May 31, 2024**
40. **Chander Vijay S Sanbhi**, DATA QUALITY AS A RELIABILITY MULTIPLIER: QUANTIFYING THE IMPACT OF STANDARDIZED WORK PROCESSES ON RELIABILITY MODEL ACCURACY, Vol. 52 No. 4 (2024): October-December 2024, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/385>
41. **Chander Vijay S Sanbhi**, DEGRADATION-AWARE SPARES OPTIMIZATION WITH WEIBULL/PHM UNCERTAINTY AND LOGISTICS DELAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/384>
42. **Chander Vijay S Sanbhi**, DIGITAL TWIN-RAM CO-SIMULATION FOR MAINTENANCE AND SPARING DECISIONS IN PROCESS FACILITIES, Vol. 52 No. 1 (2024): January-March 2024, Power System Protection and Control, ISSN-1674-3415,<https://pspac.info/index.php/dlbh/article/view/383>

43. **Chander Vijay S Sanbhi**, BAYESIAN UPDATING OF RAM MODELS FOR DYNAMIC AVAILABILITY FORECASTING UNDER REALISTIC DOWNTIME CONSTRAINTS, **Vol. 51 No. 3 (2023): July-September 2023**, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/381>
44. **Chander Vijay S Sanbhi**, HYBRID RELIABILITY MODELLING FOR ROTATING EQUIPMENT: PHYSICS-INFORMED LEARNING WITH SPARSE FAILURE DATA, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/382>
45. A. Srivastava, "Securing PHI in cloud-based healthcare systems: Challenges, frameworks, and future directions," J. Comput. Anal. Appl., vol. 33, no. 2, Aug. 2024. [Online]. Available: <https://www.eudoxuspress.com/index.php/pub/article/view/4349/3190>
46. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: <https://ijcnis.org/index.php/ijcnis/article/view/8682>
47. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4347/3189>
48. A. Srivastava, "A unified framework for secure cloud data management: Integrating governance, metadata intelligence, and privacy controls," J. Inf. Syst. Eng. Manage., vol. 9, no. 2, Apr. 2024. [Online]. Available: https://www.jisem-journal.com/download/45_A_Unified_Framework.pdf
49. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4603>