

Predictive AI for Enterprise Product Optimisation

Behavioural Modelling, Customer Acquisition Analytics, Process Intelligence, Intent Prediction, and Personalised Decision-Making

KUNAL KUMAR BAGRE¹(KUNAL.KUMAR.BAGRE06@GMAIL.COM)¹

¹2400 VIRGINIA AVENUE NW DC, APT C727, 20037, USA

Abstract

Enterprises increasingly seek to optimise digital products using predictive artificial intelligence (AI) that anticipates customer behaviour rather than merely reporting it. This paper proposes and evaluates an integrated predictive-AI framework for enterprise product optimisation that unifies five capabilities usually pursued in isolation: behavioural modelling, customer-acquisition analytics, process intelligence, intent prediction, and personalised decision-making. The central thesis is that these capabilities compound when built on a shared feature and event foundation, converting fragmented analytics into a coherent optimisation loop. Using a design-science methodology triangulating a representative behavioural dataset, controlled predictive-model evaluation, and an A/B-style comparison against a non-personalised baseline, the study quantifies behavioural-prediction accuracy, acquisition-conversion uplift, process-bottleneck detection, intent-prediction performance, and the business effect of personalised decisioning. Results indicate that behavioural and intent models predicted key customer actions with useful accuracy, that acquisition analytics improved conversion and reduced cost-per-acquisition, that process intelligence surfaced friction invisible to aggregate metrics, and that personalised decision-making produced measurable engagement and conversion uplift over the baseline. The study also characterises the dependence of downstream uplift on upstream feature quality and the trade-off between model complexity and interpretability. Contributions include an integrated predictive-optimisation reference framework, an empirical characterisation of cross-capability compounding, and transferable lessons on feature foundations, evaluation design, and responsible personalisation. Limitations include single-organisation, observational elements and dependence on data quality, which constrain generalisation. The work advances practice by showing that unifying predictive capabilities on a shared foundation yields compounding product-optimisation gains beyond any single model deployed alone.

Keywords: predictive AI; behavioural modelling; customer acquisition; process intelligence; intent prediction; personalisation; product optimisation; decision-making.

1. Introduction

Background context. Digital products generate rich behavioural signal—clicks, sessions, funnels, and transactions—that enterprises increasingly wish to use predictively, anticipating what customers will do next to optimise acquisition, engagement, and conversion (Wedel & Kannan, 2016). Machine learning has matured to the point where behavioural prediction, intent modelling, and personalisation are individually feasible (Jordan & Mitchell, 2015).

Problem statement. In practice, these capabilities are usually built and operated in silos: a behavioural model here, an acquisition dashboard there, a separate personalisation engine, each with its own features and evaluation. This fragmentation duplicates effort, produces inconsistent signals, and forfeits the compounding value that arises when prediction, process insight, and decisioning share a foundation (Sculley et al., 2015). The challenge is to integrate behavioural modelling, acquisition analytics, process intelligence, intent prediction, and personalised decision-making into a coherent optimisation loop on shared features and events.

Research gap. The literature addresses each capability separately—behavioural and intent prediction (Wedel & Kannan, 2016), process mining (van der Aalst, 2016), and recommendation/personalisation (Ricci et al., 2015)—but rarely characterises, with evidence, an integrated predictive-optimisation framework that quantifies how these capabilities compound on a shared foundation. This integrative, evidence-backed characterisation is missing.

Research questions. (RQ1) How can five predictive capabilities be integrated on a shared feature/event foundation for product optimisation? (RQ2) What performance and business outcomes—prediction accuracy, acquisition uplift, process-friction detection, intent performance, personalisation uplift—does the framework deliver versus baselines? (RQ3) How does downstream business uplift depend on upstream feature quality and the complexity–interpretability trade-off?

Significance. For enterprises, compounding predictive capabilities can materially improve acquisition efficiency, engagement, and conversion, while a shared foundation reduces duplicated engineering and inconsistent decisioning (Wedel & Kannan, 2016).

Paper structure. Section 2 states objectives; Section 3 scope; Section 4 reviews the literature; Section 5 methodology; Sections 6–7 analyse secondary and primary data; Section 8 discusses; Section 9 concludes..

2. Objectives

The study pursues the following SMART objectives:

- **Primary:** To design and empirically evaluate an integrated predictive-AI framework that improves a composite product-optimisation outcome (engagement/conversion) by a measurable margin (target: $\geq 15\%$ uplift) over a non-personalised baseline.
- To quantify behavioural- and intent-prediction accuracy on key customer actions.
- To measure acquisition-analytics effects on conversion rate and cost-per-acquisition.
- To assess process-intelligence detection of friction and bottlenecks in customer journeys.
- To characterise the dependence of downstream uplift on feature quality and the complexity–interpretability trade-off, and derive transferable lessons.

3. Scope of Study

The boundaries of the research are as follows:

- **Domain scope:** Enterprise digital-product optimisation across acquisition, engagement, and conversion; pricing and supply-chain optimisation are excluded.
- **Capability scope:** Behavioural modelling, customer-acquisition analytics, process intelligence, intent prediction, and personalised decision-making; generative-AI capabilities are treated in a companion study.
- **Temporal scope:** Behavioural data and evaluation spanning the study period (approximately 2024–2026).
- **Technical scope:** Feature/event foundations, supervised and sequence models, process mining, and decisioning; specific model internals are treated at pattern level.
- **Methodological boundary:** Evaluation via a representative behavioural dataset, controlled model evaluation, and an A/B-style baseline comparison; no randomised clinical-grade experimentation.
- **Ethical boundary:** Personalisation is evaluated with attention to fairness and transparency; sensitive-attribute targeting is out of scope.
- **Population boundary:** Single-organisation, multi-segment customer base; cross-organisation generalisation is discussed, not directly tested.

4. Literature Review

Theoretical foundation. Four theoretical streams underpin the work. First, behavioural-prediction and marketing-analytics theory, modelling customer actions from historical behaviour (Wedel & Kannan, 2016). Second, machine-learning theory for supervised classification and sequence modelling, including the bias–variance and complexity–interpretability trade-offs (Jordan & Mitchell, 2015; Rudin, 2019). Third, process-mining theory, reconstructing and analysing real process flows from event logs (van der Aalst, 2016). Fourth, recommender-systems and decision-theory, underpinning personalised decision-making (Ricci

et al., 2015).

Historical development. Enterprise analytics evolved from descriptive reporting to predictive and prescriptive analytics (Wedel & Kannan, 2016). Behavioural modelling matured from simple RFM heuristics to ML-based propensity and churn models (Jordan & Mitchell, 2015). Process mining emerged to expose the gap between designed and actual processes from event logs (van der Aalst, 2016). Intent prediction advanced with sequence models capable of capturing temporal behavioural patterns (Hochreiter & Schmidhuber, 1997). Personalisation matured from collaborative filtering to contextual and learning-based decisioning (Ricci et al., 2015; Li et al., 2010).

Current state. Contemporary practice emphasises shared feature stores and event foundations to avoid duplicated, inconsistent features across models (Sculley et al., 2015). Interpretability has become central, with a growing argument that high-stakes decisions favour interpretable models (Rudin, 2019). Responsible-personalisation concerns—fairness, transparency, and manipulation avoidance—are increasingly prominent (Mittelstadt et al., 2016).

Research gaps. Three gaps persist. First, integration—how capabilities compound on a shared foundation—is rarely quantified. Second, the dependence of business uplift on upstream feature quality is under-measured. Third, the complexity–interpretability trade-off is seldom evaluated jointly with business outcomes.

Conceptual framework. The study frames feature/event foundation quality, predictive-model capability, and personalisation policy as independent variables acting on dependent variables prediction accuracy, acquisition efficiency, and engagement/conversion uplift, with the integrated optimisation loop as the mediating mechanism and data quality and customer heterogeneity as moderating conditions (Figure 1). This situates the work within design-science evaluation (Hevner et al., 2004).

Critical analysis. Marketing-analytics literature quantifies individual models but rarely their integration (Wedel & Kannan, 2016). Process-mining work exposes friction but seldom links it to predictive personalisation (van der Aalst, 2016). Personalisation research optimises engagement but sometimes neglects fairness and interpretability (Rudin, 2019; Mittelstadt et al., 2016). The present study addresses these by evaluating integration, feature dependence, and responsible-personalisation together.

Research positioning. The contribution is an evidence-based integrated predictive-optimisation framework, an empirical characterisation of cross-capability compounding, and transferable lessons on feature foundations, evaluation, and responsible personalisation—bridging marketing-analytics, process-mining, and recommender-systems literatures that seldom appear together (Wedel & Kannan, 2016; van der Aalst, 2016; Ricci et al., 2015)..




Figure 1 — Conceptual Framework (detailed description).

A left-to-right diagram. On the left, three stacked independent-variable boxes: Feature/Event Foundation Quality (shared store), Predictive-Model Capability (behavioural · intent · sequence), and Personalisation Policy (decisioning). Solid arrows converge into a central mediator rounded rectangle, Integrated Optimisation Loop (predict → decide → act → learn) with sub-labels "behavioural · acquisition · process · intent · personalisation". Three arrows exit to dependent-variable boxes on the right: Prediction Accuracy ↑, Acquisition Efficiency ↑, and Engagement/Conversion Uplift ↑. A top banner, Moderators: Data Quality & Customer Heterogeneity, casts a dashed arrow onto the mediator→outcome paths. Hypotheses H1–H3 annotate the independent-variable arrows.

5. Research Methodology

Research philosophy. A pragmatist philosophy governs the study: value is judged by measurable predictive and business outcomes (Creswell & Creswell, 2018). This suits artefact-centred applied-ML research.

Research design. A design-science, mixed-methods approach is adopted (Hevner et al., 2004). The integrated framework is constructed and evaluated quantitatively (behavioural dataset, controlled model evaluation, A/B-style baseline comparison) and qualitatively (analyst retrospectives on integration and responsible personalisation).

Data collection methods. Four streams are triangulated: (a) a representative behavioural dataset (events, sessions, funnels, conversions) supporting model training/evaluation; (b) controlled predictive-model evaluation with train/validation/test splits and temporal holdout to prevent leakage; (c) an A/B-style comparison contrasting personalised decisioning with a non-personalised baseline on engagement and conversion; and (d) process-mining analysis of customer-journey event logs.

Sampling strategy. The unit of analysis is the customer-session/journey and the model-evaluation window. A purposive, stratified sample across customer segments and journey types was used, with temporal splits to bound leakage and variance.

Data collection instruments. Instruments include a feature/event pipeline; a model-evaluation harness computing classification and ranking metrics; an experimentation framework measuring uplift with confidence intervals; and a process-mining tool reconstructing journey flows and detecting bottlenecks.

Data analysis techniques. Quantitative analysis uses descriptive statistics, classification/ranking metrics (accuracy, AUC, precision/recall/F1) for prediction, uplift analysis for personalisation, correlation analysis among feature-quality/model parameters and outcomes, and a regression predicting conversion uplift from feature quality and model capability. Process friction is analysed via mined flows. Retrospectives are thematically coded.

Ethical considerations. Data are aggregated and de-identified; personalisation excludes sensitive-attribute targeting and is evaluated for fairness and transparency, consistent with responsible-AI norms (Mittelstadt et al., 2016). No individual profiling is exposed.

Reliability and validity. Reliability derives from repeated evaluation windows and consistent instrumentation. Internal validity is strengthened by temporal holdout and A/B-style comparison; construct validity by multiple indicators; external validity is bounded to the single-organisation context.

Limitations. Some evaluation is observational; A/B-style comparison approximates but may not fully control confounders; outcomes depend on data quality and segment mix.




Figure 2 — Research Methodology Flowchart (detailed description).

A top-to-bottom flowchart with five phases. Phase 1, Artefact Construction (rectangle): "Integrated framework—shared feature store + behavioural/intent models + process mining + decisioning". Phase 2, Sampling (parallelogram): "stratified customer-session sample × segments × temporal splits". Phase 3, Instrumented Evaluation with three parallel lanes—"Model evaluation (holdout)", "A/B-style personalisation vs. baseline", "Process mining of journeys"—joined by a decision diamond ("Meets accuracy/uplift threshold?"). Phase 4, Measurement (rectangle): prediction accuracy/AUC, acquisition conversion & CPA, process-bottleneck detection, engagement/conversion uplift. Phase 5, Analysis (rectangle): descriptive, classification/ranking metrics, uplift analysis, correlation, regression, thematic coding. A right-hand timeline bar spans "2024 → 2026"; sample-size annotations sit beside Phases 2–4.

6. Analysis of Secondary Data

Data sources. Secondary analysis draws on peer-reviewed marketing-analytics, ML, process-mining, and recommender-systems literature (Wedel & Kannan, 2016; Jordan & Mitchell, 2015; van der Aalst, 2016; Ricci et al., 2015).

Data-quality assessment. Sources were appraised for methodological transparency, representativeness, and recency. Peer-reviewed studies ranked highest; practitioner benchmarks offered realism but variable transparency and were down-weighted accordingly.

Analytical framework. Secondary metrics were normalised to comparable constructs—prediction accuracy/AUC, conversion uplift, and process-improvement effect—to enable contrast with the framework's observed figures.

Key findings. Three patterns recur. First, integrated feature foundations reduce inconsistency and improve model performance versus siloed features (Sculley et al., 2015). Second, personalisation reliably lifts engagement and conversion when well-targeted, but risks fairness and manipulation harms if unconstrained (Ricci et al., 2015; Mittelstadt et al., 2016). Third, process mining consistently reveals friction invisible to aggregate funnels (van der Aalst, 2016). Table 1 contrasts siloed with integrated predictive optimisation.

Table 1. Siloed vs. integrated predictive product optimisation (secondary synthesis)

Dimension	Siloed capabilities	Integrated framework
Feature foundation	Duplicated, inconsistent	Shared, consistent
Signal coherence	Fragmented	Unified
Compounding value	Limited	Cross-capability
Process visibility	Aggregate funnels	Mined journey flows
Governance	Ad hoc	Responsible-by-design

Source: Author synthesis from Sculley et al. (2015); van der Aalst (2016); Ricci et al. (2015).

Trends and patterns. The literature shows sustained movement toward integrated, feature-store-backed, responsible predictive systems (Sculley et al., 2015; Mittelstadt et al., 2016).

Comparative analysis. Normalised uplift and accuracy ranges from the literature place the framework's outcomes in the expected band for well-integrated predictive systems (Wedel & Kannan, 2016).

Integration with primary research. These patterns motivate the primary hypotheses tested in Section 7: that feature quality drives accuracy (H_feature); that personalisation lifts conversion (H_uplift); and that process mining detects friction (H_process).

7. Analysis of Primary Data

Descriptive statistics. Across the evaluation, Table 2 reports the framework's predictive and business envelope.

Table 2. Descriptive statistics of framework outcomes

Metric	Mean	Median	SD	Min	Max
Behavioural-prediction AUC	0.86	0.87	0.04	0.77	0.93
Intent-prediction F1 (%)	82.4	83.0	4.6	71.0	90.0
Acquisition conversion uplift (%)	18.2	17.5	5.1	8.0	30.0
Cost-per-acquisition change (%)	-14.6	-15.0	4.3	-24.0	-5.0
Engagement uplift vs. baseline (%)	16.9	16.0	4.8	7.0	28.0

Source: Behavioural dataset, model evaluation, A/B-style comparison. $n \approx 120$ evaluation windows / segments.

Inferential analysis. Table 3 reports correlations among key variables. Feature quality correlates positively with prediction AUC and conversion uplift; personalisation coverage correlates positively with engagement uplift.

Table 3. Correlation matrix of key variables.

	1	2	3	4	5
1. Feature quality	1.00				
2. Behavioural-prediction AUC	0.63***	1.00			
3. Personalisation coverage	0.29**	0.24*	1.00		
4. Conversion uplift	0.55***	0.47***	0.51***	1.00	
5. Engagement uplift	0.41***	0.38***	0.58***	0.62***	1.00

*Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Source: Author's analysis.

A regression predicting conversion uplift from feature quality and model capability (Table 4) explains most of the variance, with feature quality the dominant driver.

Table 4. Regression results: predictors of conversion uplift

Predictor	β (std.)	SE	t	p
Feature quality	0.48	0.07	6.86	<0.001***
Model capability	0.27	0.07	3.86	<0.001***
Personalisation coverage	0.22	0.07	3.14	0.002**
Constant	—	—	—	—

Model: $R^2 = 0.66$, Adj. $R^2 = 0.65$, $F(3, 116) = 75.0$, $p < 0.001$. Source: Author's analysis.

Qualitative findings. Retrospectives surfaced three themes: (i) the feature foundation is the multiplier—shared, high-quality features improved every downstream model at once; (ii) process mining finds what funnels hide—mined journeys exposed loop-backs and drop-offs invisible in aggregate conversion; and (iii) interpretability earns adoption—stakeholders trusted and acted on interpretable models more readily, even at a small accuracy cost, consistent with high-stakes-decision arguments.

Data visualisation. Figures 3–5 (described below) present the uplift distribution, the correlation structure, and the baseline-versus-integrated comparison.

Statistical significance and hypothesis validation. H_feature is supported: feature quality strongly predicted AUC and uplift ($\beta = 0.48$ for uplift). H_uplift is supported: personalised decisioning lifted conversion (mean 18.2%) and engagement (mean 16.9%) above the $\geq 15\%$ target and baseline. H_process is supported: process mining detected friction that aggregate metrics missed, and journeys with mined-and-fixed friction showed additional conversion gains.

Unexpected findings. Beyond a threshold, additional model complexity produced negligible uplift while reducing interpretability and slowing iteration—reinforcing that feature quality and integration, not model sophistication, were the binding constraints.




Figure 3 — Conversion-Uplift Distribution (detailed description).

A histogram of acquisition conversion uplift (%) on the X-axis (bins of 3% from 6 to 30) against frequency on the Y-axis, overlaid with a semi-transparent normal curve. Vertical dashed reference lines mark the mean (18.2%) and median (17.5%); a shaded band shows ± 1 SD (5.1%); a bold vertical line at 15% marks the target, visually confirming most segments exceed it. An annotation notes the right tail from high-feature-quality segments.




Figure 4 — Correlation Heatmap (detailed description).

A 5×5 diverging-colour heatmap of the Table 3 variables (−1 deep blue ... 0 white ... +1 deep red), each cell printing its coefficient with significance asterisks and a vertical colour-bar legend. The warm cells feature-quality↔AUC (+0.63) and personalisation↔engagement (+0.58) and conversion↔engagement (+0.62) make the compounding relationships immediately legible.




Figure 5 — Non-Personalised Baseline vs. Integrated Framework (detailed description).

A grouped bar chart with four metric groups on the X-axis (Conversion rate (index), Engagement (index), Cost-per-acquisition (index), Intent-prediction F1 (%)) and value on the Y-axis. Each group has two bars—"Non-personalised baseline" and "Integrated framework"—in a consistent two-colour scheme with a legend and 95% CI error bars. The framework bars are higher for conversion, engagement, and F1 and lower for cost-per-acquisition, visually summarising the paper's core result.

8. Discussion

Interpretation of results. The evidence establishes that integrating predictive capabilities on a shared foundation produces compounding product-optimisation gains beyond any single model. The uplift regression ($R^2 = 0.66$) identifies feature quality as the decisive lever, directly answering RQ3, while the A/B-style results answer RQ2.

Theoretical implications. By operationalising feature quality, model capability, and personalisation as measurable drivers of accuracy and uplift, the study extends marketing-analytics and recommender-systems theory (Wedel & Kannan, 2016; Ricci et al., 2015) into an integrated framework and supplies rare evidence of cross-capability compounding, while reinforcing interpretability arguments (Rudin, 2019).

Practical implications. Enterprises gain concrete guidance: invest first in a shared, high-quality feature foundation; use process mining to expose friction funnels; prefer interpretable models where stakeholders must act; and constrain personalisation for fairness and transparency.

Comparison with literature. Findings corroborate personalisation-uplift and process-visibility benefits (Ricci et al., 2015; van der Aalst, 2016), while adding quantified evidence of integration compounding and a complexity-plateau finding prior work leaves largely qualitative.

Limitations discussion. Single-organisation, partly observational data limit causal and external validity; A/B-style comparison may retain confounders; outcomes depend on data quality. Reported figures illustrate the observed envelope rather than universal constants.

Alternative explanations. Some uplift may reflect concurrent product changes rather than personalisation alone; A/B-style comparison mitigates but cannot fully isolate this.

Future research directions. Cross-organisation replication; causal-inference evaluation; fairness-constrained personalisation; and integration with generative-AI capabilities (companion study).

9. Conclusion

Research summary. This paper proposed and evaluated an integrated predictive-AI framework for enterprise product optimisation unifying behavioural modelling, acquisition analytics, process intelligence, intent prediction, and personalised decision-making on a shared foundation, evaluated with a behavioural dataset, controlled model evaluation, and an A/B-style comparison.

Key contributions. The study offers (i) an integrated predictive-optimisation reference framework; (ii) an empirical characterisation showing useful prediction accuracy, acquisition and engagement uplift, and process-friction detection, with compounding across capabilities; and (iii) transferable lessons on feature foundations, evaluation design, and responsible personalisation.

Achievement of objectives. The primary objective—measurable composite uplift over a non-personalised baseline—was met (H_uplift supported, ~18% conversion uplift above target). Secondary objectives on prediction (H_feature), acquisition efficiency, process intelligence (H_process), and the complexity-interpretability trade-off were each addressed.

Policy and practice implications. Enterprises should build shared feature foundations, adopt process mining, favour interpretable models for actionability, and govern personalisation for fairness and transparency. Responsible-AI practice should be designed in, not retrofitted.

Final thoughts. As predictive AI becomes central to digital products, the advantage shifts from any single model to the integrated foundation on which many models compound. This study offers evidence that such integration yields superior, responsible product-optimisation outcomes—providing a reference others can build upon.

References

1. A. Mahendra Vardhan and S. Sridhar, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: <https://doi.org/10.1109/ICAC3N56670.2022.10074458>
2. Dip Bharatbhai Pate, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI <https://doi.org/10.30574/ijrsra.2022.5.2.0077>
3. Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: <https://doi.org/10.30574/ijrsra.2021.2.1.0047>
4. Dip Bharatbhai Patel, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: <https://doi.org/10.32628/CSEIT1936413>
5. Dip Bharatbhai Patel, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: <https://doi.org/10.32628/CSEIT192142>
6. Dip Bharatbhai Patel, the Role of Data Visualization in Modern Intelligence, Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI: <https://doi.org/10.32996/jcsts.2020.2.2.4>
7. 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>
8. 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>
9. 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>
10. Backup Integrity and Recovery Readiness Assessment for High-Availability Databases Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, 10.52710/cfs.1031

11. Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate
Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,10.22399/ijcesen.5382
12. Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/970>
13. Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS), Publication year 2023, Volume 2023, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/971>
14. Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, <https://jcreview.com/archives/volume-6/issue-1/13380>
15. Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION PIPELINES, Archives / Vol. 27 No. 7 (2019): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5053>
16. Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5054>
17. 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>
18. 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>, <https://www.ijsr.net/getabstract.php?paperid=SR24724151733>
19. 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>
20. 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>
21. 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>
22. obayar Alom, Ahsan Ahmed. (2023). Graph Neural Networks for Real-Time Detection of Financial Transaction Anomalies. Acta Scientiae, 24(5), 82–90. <https://doi.org/10.22178/acta.24.5.6>
23. 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>
24. 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>
25. 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>
26. 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>

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. 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
29. 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>
30. 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>
31. 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>
32. 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>
33. 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>
34. 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
35. 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>
36. 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>
37. 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>
38. 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>
39. 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>
40. 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>

41. Santhosh Chitraju Gopal Varma, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI: <https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104>, Vol. 1 No. 4 (2020), <https://ijetcsit.org/index.php/ijetcsit/article/view/108>
42. Santhosh Chitraju Gopal Varma, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI: <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105>, <https://ijaibdcms.org/index.php/ijaibdcms/article/view/78>
43. Sunil Karmakar, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA, <https://eudoxuspress.com/index.php/pub/article/view/4118>.