



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

The Role of AI in Startup Resource Optimization: A Content-Analytic Study Using Chi-Square, ANOVA and Linear Regression

C. Lakshmi Sree Varshitha¹, V. Narendra²

¹. Research Scholar, BBA Department, *Koneru Lakshmaiah Education Foundation*, Vaddeswaram, Andhra Pradesh, India.

². Assistant Professor, BBA Department, *Koneru Lakshmaiah Education Foundation*, Vaddeswaram, Andhra Pradesh, India.

DOI : <https://doi.org/10.55248/gengpi.07.1026.3254>

ABSTRACT

Startups operate using a limited amount of financial, human, temporal and technological resources. The effective application of these scarce resources determines the survival and growth prospects of a startup. Artificial Intelligence, referred to as AI is increasingly being leveraged to perform tasks look at business data predict demand assist in decision-making and enhance customer service. This study examines the impact of Artificial Intelligence on the optimization of a startup's resources. The research methodology uses a set of ten articles published in 2024-2026 in Bloomberg, Fortune and the Indianapolis Business Journal periodicals. For each article we identified the resource that was emphasized assessed the level of Artificial Intelligence utilization and determined the presence of five resource-optimization areas. These areas include cost efficiency, time efficiency, employee productivity, decision-making and operational efficiency. We applied a set of statistical instruments including chi-square tests, one-way ANOVA and simple linear regression to analyse the data. We present the results of our findings using bar graphs.

Keywords: Artificial Intelligence, startups, resource optimization, cost efficiency, chi-square, ANOVA regression

INTRODUCTION

AI is gaining a prominent place in the business world. Businesses are using it to carry out tasks, analyze data, understand customers and aid decision-making, and it can assist companies in streamlining processes and reducing manual workload.

Startups usually operate on limited finances. They are constrained in finance, staff, time, technology and operating capacity, which makes it difficult to manage resources carefully. Using resources efficiently allows a startup to save costs, raise productivity and monitor its activities more closely.

AI can assist in many areas of resource management. It can automate repetitive tasks and free employees for higher-value work. AI-powered data analysis unveils consumer preferences and market trends, and AI can also aid demand forecasting, inventory planning, marketing, customer service and business decision-making.

Adopting AI is not always easy for startups. Tools can be expensive, technical expertise may be lacking, employees need training, and founders may not know which tools to choose. It is therefore important to understand the actual impact of AI on resource optimization. This study asks whether the use of AI helps startups use their available resources more efficiently in cost management, time, employee productivity, decision-making and operational efficiency.

LITERATURE REVIEW

How is Business Being Influenced by AI? The increase in adoption of AI has already impacted business by increasing efficiency, risk assessment, customer service, investment analysis and decision-making. The sources analysed were published from 2024 to 2026 and have all seen some form of implementation of AI in financial institutions, fintechs, startups or other. These ten sources represent the dataset analysed in Section 7. Weiss (2025) reported about Cipher Owl, an AI-focused crypto compliance startup that has attracted \$15m to improve transaction monitoring and fraud detection. Although the startup has only 8 staff it was able to attract accounts with Coinbase and OKX, and the firm leverages generative AI to generate reports on flagged transactions; indicating that AI can reduce time needed for transaction review and empower a tiny staff to manage a multitude of digital transactions.

Gordon (2025) explained how large banks and fintech startups use AI. Big players are slow as they have reputational, regulatory and operational risk issues. Startups may be more risk tolerant and more likely to make riskier bets. This article described "AgentFi", in which an AI agent may one day engage in financial transactions on your behalf, suggesting that AI could augment decision-making for finance or even automate processes entirely. Lee (2024) Upstage, which raised \$72 million in Series B funding to develop enterprise-AI products, including document-processing engines and the

Solar large language model. The growth illustrates how AIs can boost enterprise productivity, by automating the work of processing information. Lee (2025) investigated how Daiwa Securities formed a collaboration with Sakana AI to develop an AI-driven platform for retail investment services that supports staff preparing proposals, market analysis and portfolio recommendations. This is an example of how wealth management professionals use AI to handle and personalise the service they provide.

In 2026, Lee stated that "AI chip startup Furiosa AI announced a US\$500 million funding round before IPO". This article's bibliographic information was obtained and used for descriptive coding only.

AI infrastructure acquisition by Nvidia (Schuetz 2024) Nvidia announced an acquisition of Run: Ai, a company that develops AI infrastructure-management software. Nvidia's 2024 acquisition of the AI infrastructure company, Run: Ai, is estimated to be worth \$700 million and highlights the strategic importance of AI infrastructure and handling r...

Research gap

The previous research on the topic of AI, digital transformation, business performance, and technology within organizations does not focus on the advantage of AI and how it can help a startup to make the most of the scarce resource. A startup has tighter constraints than large organizations on money, people, time, and technology; it is, therefore, essential to examine the role AI plays to support them with being efficient with those resources. This paper is filling that gap by studying AI, and its influence on cost efficiency, time, employee productivity, decision-making, and operational efficiency.

OBJECTIVES OF THE STUDY

1. To learn the ways in which AI can be leveraged to optimize a startup's resources.
2. To investigate how far start-up and affiliated companies use the AI.
3. To identify ways in which AI can help decision-making in business.
4. To determine the significant challenges of AI deployment.
5. To examine the relationship between AI adoption (integration) and resource utilization.

HYPOTHESES

The hypotheses of the initial proposal are included and formalised for the dataset. The "AI adoption" variable is the 'levels of AI integration' coded for each article (7.3.), and each resource-optimisation variable is coded as 0/1 if the article provides evidence for it.

Table 1: Hypotheses

Hypothesis	Dimension	Statement (null / alternative)
H ₀₁ / H ₁₁	Cost efficiency	No significant relationship / a significant relationship between AI integration and evidence of cost efficiency.
H ₀₂ / H ₁₂	Time efficiency	No significant relationship / a significant relationship between AI integration and evidence of time efficiency.
H ₀₃ / H ₁₃	Employee productivity	No significant relationship / a significant relationship between AI integration and evidence of employee productivity.
H ₀₄ / H ₁₄	Operational efficiency	No significant relationship / a significant relationship between AI integration and evidence of operational efficiency.
H ₀₅ / H ₁₅	Overall resource optimization	No significant relationship / a significant relationship between AI integration and the overall resource-optimization score.

RESEARCH DESIGN AND METHODOLOGY

The design employed in this study is of a quantitative, descriptive and cross-sectional nature. Since the data provided by the database is the abstracts of published papers rather than questionnaire answers, the designed method of data collection with Likert-style answers was not used and was substituted for a quantitative content analysis of secondary data.

Data source and sampling

The data consist of the ten articles (covered in the file 'AIFinanceArticles_Summary.xlsx') - six from Bloomberg, three from Fortune, and one from the Indianapolis Business Journal (all published in 2024–2026). All ten articles have been analysed (a census of the dataset, purposive and non-probability). The unit of analysis is the article. As all the articles are journalistic reports and not empirical research, none have a sample design of their own.


Variables and coding

The title, variables, results and novelty fields of the articles were coded. The researcher was the sole coder, and the coding scheme is below.

DATA ANALYSIS AND RESULTS

Descriptive analysis


Among the ten articles, 2 articles were published in 2024, four articles in 2025, and four articles in 2026 (see Figure 1). Deal values for eight articles were provided and ranged from \$15 million to \$700 million (mean = \$271.8 million, SD = 292.7, median = \$121 million). The mean resource-optimisation score for the nine scorable articles was 2.00 (out of 5) (SD = 1.12).



Articles in the dataset by publication year

Source: Author's computation from the study dataset.

Figure 4 displays the frequency of each resource dimension being evidenced. Operational efficiency was present in six of the nine articles (67%), employee productivity in four (44%), time efficiency and decision-making in three each (33%), and cost efficiency in two (22%).




Chi-square tests of independence

Test 1: primary resource focus and publication time frame. The association was not significant, $(2, N = 10) = 3.056, p = 0.217$, Cramér's $V = 0.553$. The effect size is large, but with all expected counts less than 5 the p-value is not valid. All three decision and workforce articles were published in 2024-25, while two of three AI-infrastructure articles were published in 2026

Table 3: Primary resource focus × publication period (observed counts)

Primary resource focus	2024–25	2026	Total
AI infrastructure	1	2	3
Decision & workforce support	3	0	3
Operational automation	2	2	4
Total	6	4	10

Chi-square = 3.056, df = 2, p = .217, Cramér's $V = .553$. All expected counts < 5.




Test 2 (H_{05}): AI-integration level and resource-optimization evidence band. The association was not significant, $\chi^2(2, N = 9) = 3.000, p = 0.223, V = 0.577$. Level 3 (AI-native) articles were more often in the high-evidence band (2 of 3) than Level 2 articles (0 of 3), but the pattern is not monotonic and cannot be distinguished from chance (Figure 3).

Table 4: AI-integration level × evidence band (observed counts)

AI-integration level	High (3–5)	Low (0–2)	Total
Level 1 (infrastructure / strategy)	1	2	3
Level 2 (deployed in workflows)	0	3	3
Level 3 (AI-native / core layer)	2	1	3
Total	3	6	9

Chi-square = 3.000, df = 2, p = .223, Cramér's V = .577. All expected counts < 5.



Test 3: publisher group and article type. All Bloomberg articles were deals reports, 2 of 4 from Fortune/IBJ were analyses, (1, N = 10) = 3.75, $p = 0.053$, $V = 0.612$. It is close but not significant; Fisher's exact test results in $p = 0.133$. This is not a test of a hypothesis, but rather a descriptive test of the makeup of the dataset.


One-way ANOVA

ANOVA 1 (H_{05}): resource-optimization score by AI-integration level. Mean scores were 1.67 (Level 1), 1.67 (Level 2) and 2.67 (Level 3). The differences were not significant, $F(2, 6) = 0.75$, $p = 0.512$, $\eta^2 = 0.20$; the Kruskal–Walli's test agreed, $H = 1.21$, $p = 0.546$. The null hypothesis H_{05} is therefore retained (Figure 5).

Table 5: ANOVA – resource-optimization score by AI-integration level

Source	SS	df	MS	F	p
Between groups	2.000	2	1.000	0.750	0.512
Within groups	8.000	6	1.333		
Total	10.000	8			

$n = 3$ articles per level. $\eta^2 = .200$.




ANOVA 2: deal value by publication year. Mean deal values were \$386.0m (2024), \$24.5m (2025) and \$338.3m (2026). There was no difference, $F(2, 5) = 0.96$, $p = 0.445$, $\eta^2 = 0.28$ (Kruskal-Wallis $p = 0.127$), with the 2025 articles reporting smaller deals, however with only 2 in this group no trend can be suggested (Figure 6).

Table 6: ANOVA – deal value (USD million) by publication year

Source	SS	df	MS	F	p
Between groups	166060.3	2	83030.1	0.957	0.445
Within groups	433669.3	5	86733.9		
Total	599729.5	7			

$n = 2$ (2024), 2 (2025), 4 (2026). $\eta^2 = .277$.



ANOVA 3: deal value by primary resource focus. The only significant test was dealing value among different primary resource focus, $F(2, 5) = 40.23$, $p < .001$, $\eta^2 = 0.94$. The mean deal values were: \$614.3 million for AI-infrastructure deals, \$74.3 million for operational-automation deals, and \$34.0 million for the single decision-and-workforce deal (Figure 7). The wide gap between AI-infrastructure deals and the others should be interpreted with

caution: one group has only one observation, no post-hoc test was performed, and the large gap is due to the size of the infrastructure transactions (\$1.09 billion IPO stage round, two acquisitions), not any particular consequence of AI on start-up resource consumption.

Table 7: ANOVA – deal value (USD million) by primary resource focus

Source	SS	df	MS	F	p
Between groups	564640.1	2	282320.0	40.229	.001
Within groups	35089.4	5	7017.9		
Total	599729.5	7			

$n = 3$ (AI infrastructure), 4 (operational automation), 1 (decision & workforce support). $\eta^2 = .941$; $p = .0008$.




Table 8: ANOVA – resource dimensions by AI-integration level

Dimension	Mean L1	Mean L2	Mean L3	F(2, 6)	p	η^2
Cost efficiency (H ₀₁)	0.33	0.33	0.00	0.50	0.630	0.14
Time efficiency (H ₀₂)	0.33	0.00	0.67	1.50	0.296	0.33
Employee productivity (H ₀₃)	0.00	0.67	0.67	2.00	0.216	0.40
Decision-making	0.33	0.33	0.33	0.00	1.000	0.00
Operational efficiency (H ₀₄)	0.67	0.33	1.00	1.50	0.296	0.33

Means are proportions of articles evidencing the dimension ($n = 3$ per level).

4 Linear regressions


Regression 1 (H): resource-optimization score regressed on AI-integration level, $N=9$. (unsignificant positive slope: Score = $1.000 + 0.500$ Level; $b = 0.500$, $SE = 0.450$, $t(7) = 1.111$, $p = 0.303$, 95% CI [-0.56, 1.56]). 15.0% variance explained (adjusted $R = 0.029$, $R = 0.150$), $F(1, 7) = 1.235$, $p = 0.303$. Residuals satisfied normality (Shapiro-Wilk $p = 0.553$) and independence (Durbin-Watson = 2.09). For each subsequent transition in integration level, the resource-optimisation score is estimated to be half a point higher than for the previous level, but the 95% confidence interval of this value crosses zero, so H is kept (see Figure 8).

Table 9: Regression 1 – resource-optimization score on AI-integration level

Predictor	b	SE	t	p	95% CI
-----------	---	----	---	---	--------

Predictor	b	SE	t	p	95% CI
Constant	1.000	0.972	1.029	0.338	[-1.30, 3.30]
AI-integration level	0.500	0.450	1.111	0.303	[-0.56, 1.56]

$R = 0.387$, $R^2 = 0.150$, adjusted $R^2 = 0.029$, $F(1, 7) = 1.235$, $p = 0.303$, $n = 9$. Standardised $\beta = 0.387$.



Regression 2: deal value (USD million) was regressed on publication year ($n = 8$). A little bit of a flat line here... $b = 6.82$, $SE = 134.78$, $t(6) = 0.051$, $p = 0.961$, $R = 0.0004$. As publication year does not predict deal value, we have some slight deviations from normality (Shapiro-Wilk $p = 0.032$) and positive autocorrelation (Durbin-Watson = 1.08) for the residuals, so we show this for completeness only (Figures 9 and 10).

Table 10: Regression 2 – deal value (USD million) on publication year

Predictor	b	SE	t	p
Constant	-13536.8	272964.9	-0.050	0.962
Publication year	6.82	134.78	0.051	0.961

$R^2 = 0.0004$, adjusted $R^2 = -0.166$, $F(1, 6) = 0.003$, $p = 0.961$, $n = 8$.





Figure 9: Regression 2 – actual vs predicted deal / funding value by article

Source: Author's computation from the study dataset.

Figure 10: Variance explained (R^2) by the two regression mode

5 Summary of hypothesis testing

Table 11: Hypothesis decisions

Hypothesis	Test used	Result	Decision
H ₀₁ Cost efficiency	ANOVA 4	F = 0.50, p = 0.630	Not rejected
H ₀₂ Time efficiency	ANOVA 5	F = 1.50, p = 0.296	Not rejected
H ₀₃ Employee productivity	ANOVA 6	F = 2.00, p = 0.216	Not rejected
H ₀₄ Operational efficiency	ANOVA 8	F = 1.50, p = 0.296	Not rejected
H ₀₅ Overall optimization	ANOVA 1, Regression 1, Chi-square 2	p = 0.512; 0.303; 0.223	Not rejected

DISCUSSION

None of the five null hypotheses could be rejected; there is no statistically significant association between the quantity of resource-optimisation evidence and the degree of AI integration. This would not constitute evidence of no association; the tests are extremely underpowered at nine scorable articles per experiment, as is demonstrated by the regression slope itself, which of 0.5 points per degree of AI integration has a confidence interval of 0.56-1.56. A number of medium-to-large effect sizes (.14 to .40 for most of the tests, V of 0.55-0.61 for the chi-square tests) lean in the direction of the hypothesis as implied by the proposal; however, given the small sample size, these cannot be deemed reliable.

These descriptive findings are more interesting. Efficiency was evidenced more frequently than effectiveness, followed by employee productivity, aligning with the conceptual expectation that automation is the primary vehicle for AI to create efficiency in startup resources. Cost efficiency was evidenced less frequently (2 of 9 articles), although the most pronounced numerical example in the data set is a cost one: Orr (2025) found that the cost of building a minimum viable product that once would have taken about two million dollars is now down to about two hundred thousand dollars (a 90 percent cost reduction), and the number of employees employed by the startup was reduced from eight to three to five employees. ProBook's clients experienced fewer dispatchers and higher average revenue per job (Lazarus, 2026), and Cipher Owl's largest clients have eight employees (Weiss, 2025).

That the ANOVA on deal value by focus is also large simply indicates AI-infrastructure deals are orders of magnitude bigger than firms working with AI on operations. It bears on the implementation-issues hypothesis because the AI infrastructure cost and compute itself is sufficient to draw hundreds of millions of dollar acquisitions (Schuetz, 2024; Bass, 2026), aligning with the concern that AI tools can be expensive for resource-constrained startups. Gordon (2025) notes large institutions being cautious due to hallucination, deepfake, and agent-specific risks--additional implementation issues.

MAJOR FINDINGS

Aim 1 (using AI to make better use of resources): the articles demonstrate AI primarily in terms of working efficiency (67%), employee productivity (44%), time efficiency and improved decision making (33% each) and cost efficiency (22%).

Objective 2 (level of adoption): six of the nine articles that received a score describe AI being used within a workflow (Level 2) or as the flexible layer on which the business operates (Level 3), with the number of AI-oriented startups in the dataset increasing from two articles in 2024 to four articles in 2025 and 2026.

Objective 3 (business decisions): decision support appears in three of nine articles, in wealth-management advice (Daiwa), dispatch scheduling (ProBook) and AI-agent finance (Gordon, 2025).

Objective 4 (implementation challenges): the challenges are the expensive AI compute and infrastructure, reputational and regulatory risk, AI-specific risks such as hallucination and deepfakes, and slow contracting between startups and large institutions.

-Null Research Hypotheses-I V (number of nulls retained): The link is in a positive direction ($b = 0.50$) but is not significant ($p = .303$). All five null hypotheses are retained.

SUGGESTIONS AND IMPLICATIONS

Startups should begin with operational automation (scheduling, reporting, document management), where the proof point of benefit is clearest, before branching out into more sophisticated AI-agent initiatives.

Cost of AI tools versus headcount and development cost-AI infrastructure can also be costly, so founders need to balance the costs of the AI tools versus the savings they provide compared to the headcount.

Investors and policymakers can support startups by fostering shared compute access and low-cost AI tooling.

Future studies should gather primary data from startup founders and employees using the proposed questionnaire (5-point Likert scale), and verify if the tests performed in the present study work on that data.

LIMITATIONS

There were only ten articles on which the dataset was based (nine had full text) so all tests are likely to be underpowered and results exploratory. Articles are news reports, not scientific papers, they may be biased towards positive outcomes, and many focus on AI firms not using AI to save resources.

Coding was performed by one coder, and intercoder reliability was not calculated, and the AI-integration dimension flags involved subjective judgements.

All chi-squares have expected counts less than 5, so asymptotic p-values not valid. ANOVA 3 has a group with only one point. Residuals of Regression 2 not normally distributed.

Article A5 contributed bibliographic details only.

CONCLUSION

This thesis used chi-square tests, one-way ANOVA and linear regressions of ten coded articles to investigate whether AI has a positive effect on startups' resource optimization in these articles, tested to the proposition. AI was most common in evidence as an efficiency/employee productiveness enhancer, where the clearest example was quantified savings in startup development and staffing costs. Statistically speaking, however, the degree of AI use was not a significant predictor of resource-optimization evidence, with the only significant finding, larger deal sizes for AI consistent firms, relating to cost of AI rather than benefits.

This thesis underpins the direction of effect in the proposition, but not yet its statistical claims.

A questionnaire-based research design with an appropriate sample of startup founders and employees is the obvious follow-up.

REFERENCES

Bass, D. (2026). Nebius agrees to buy startup that makes AI run faster, cheaper. Bloomberg.com.

Gordon, N. (2025). Big banks may 'afford to be a bit behind the curve' on AI, and leave riskier bets to smaller startups. Fortune.com.

Lazarus, L. M. (2026). Exclusive: A founder who went from pressure washing to Wharton just raised \$40 million to put AI to work for plumbers, electricians, and HVAC crews. Fortune.com.

Lee, M. J. (2025). Daiwa employs AI startup Sakana to create application for retail customers. Wwww. Bloomberg.com.

Lee, Y. (2024). AI startup Upstage bags \$72 million for US and Asian growth. Bloomberg.com.

Lee, Y. (2026). AI chip startup Furiosa AI readies for \$500 million round ahead of IPO. Bloomberg, L. (2026).

Orr, S. (2025). AI tools help startups slash their early expenses. Indianapolis Business Journal.

Schuetz, M. (2024). Nvidia to buy Israeli AI software firm Run: Ai. Bloomberg.com.

Sun, Y. (2026). AI startup Dwelly raises \$170m for real estate rollup play. Bloomberg.com.

Weiss, B. (2025). Exclusive: Coinbase and Cruise alums raise \$15 million for crypto compliance startup Cipher Owl. Fortune.com

APPENDIX: CODED DATASET

ID	Article	Year	Publisher	Deal (\$m)	Focus	Level	Cost	Time	Emp.	Dec.	Ops	Score
A1	Cipher Owl	2025	Fortune	15	Operational	3	0	1	1	0	1	3
A2	Big banks & AI	2025	Fortune	–	Decision	1	0	0	0	1	0	1
A3	Upstage	2024	Bloomberg	72	Operational	2	0	0	0	0	1	1
A4	Daiwa–Sakana	2025	Bloomberg	34	Decision	2	0	0	1	1	0	2
A5	Furiosa AI	2026	Bloomberg	500	Infrastructure	–	–	–	–	–	–	–
A6	Nvidia–Run: Ai	2024	Bloomberg	700	Infrastructure	1	0	0	0	0	1	1
A7	Nebius–Eigen	2026	Bloomberg	643	Infrastructure	1	1	1	0	0	1	3
A8	ProBook	2026	Fortune	40	Operational	3	0	1	1	1	1	4
A9	Dwelly	2026	Bloomberg	170	Operational	3	0	0	0	0	1	1
A10	IBJ startups	2025	IBJ	–	Decision	2	1	0	1	0	0	2

Cost, Time, Emp. (employee productivity), Dec. (decision-making), Ops (operational efficiency) are 0/1 flags. – = not available.