CINEFORUM
ISSN: 0009-7039
Vol. 66. No. S2, 2026 (Special Issue S2)

Evaluating the Impact of AI-Augmented Blended Learning on Sales Performance: a Pre–Post Study

Deepak Verma[1]*, Dr. Mani Tyagi[2], Himanshu Saxena[3]

[1]Research Scholar Krishna Institute of Engineering & Technology (KIET), Ghaziabad, Delhi-NCR, Uttar Pradesh, India
[2] Professor Krishna Institute of Engineering & Technology (KIET), Ghaziabad, Delhi-NCR, Uttar Pradesh, India
[3] Associate Professor Krishna Institute of Engineering & Technology (KIET), Ghaziabad, Delhi-NCR, Uttar Pradesh, India
***Corresponding Author**
Deepak Verma
Research Scholar Krishna Institute of Engineering & Technology (KIET), Ghaziabad, Delhi-NCR, Uttar Pradesh, India
Email ID:-deepak.26sphd0075@kiet.edu

ABSTRACT
Purpose – Selling a used car usually comes down to a single conversation between the rep and the buyer, and so the rep's training carries a lot of weight. AI-augmented training is the most recent attempt at building this skill, and also the one with the least evidence behind it. The question here is simple: when an on-demand, AI-driven conversational module was added to a human-trainer model held back by location and trainer availability, did outlet-level sales performance move?
Design – The study followed a single-group interrupted time series design at three outlets of a large, organised used-car retailer in Delhi NCR. Twelve months before the AI module went live (2024) were set against the twelve months after it (2025), giving 72 outlet-months in all. The analysis combined segmented ITS regression with a two-way ANOVA (outlet × period), read through the Job Demands-Resources model, and every step was run in Python.
Findings – At the point of launch, neither net deliveries nor tokens jumped in any significant way. What changed was the direction of travel afterwards, with a significant and lasting upward trend (net deliveries +23.5/month, tokens +47.1/month, both $p < .001$) and a 38.2% drop in return rate over the same period. The three outlets did not differ significantly in how much they improved ($F(2,66) = 1.87–2.36$, $p > .10$). Since only three outlet clusters stand behind this test, it is best read as a modestly powered result and not as firm evidence that the impact was uniform.
Originality – Among the peer-reviewed studies on AI-supported training examined for this paper, automotive and used-car sales settings barely appear, and few of them put an AI tool up against real transactional outcomes. This study offers one such evaluation from India and represents an early stage of a broader research programme.
Keywords: AI-augmented training; Sales performance; Job Demands-Resources model; Interrupted time series analysis; Blended learning; Used-car retail

1. INTRODUCTION

A good share of used-car sales are decided in a single conversation between the rep and the customer. How well that rep has been trained therefore shows up quite directly in the numbers. Keeping training consistent becomes hard once outlets are scattered over a large area, simply because a human trainer can only be in one place on any given day. AI-augmented training has drawn attention as one way out of this bind. Some organisations use it to replace trainers altogether; others, including the one studied here, add it on as a supporting layer. That is broadly how AI has entered HR development too: organisations have taken it up well before anyone checked, with much care, what it actually delivers (Pereira et al., 2023). Sales training shows the gap more sharply. Used-car retail shows it most of all, because what little evidence exists comes from trainees describing their training, not from the cars they later sold.

2. THEORETICAL FRAMEWORK

Much of what a sales rep knows about selling is learnt after joining, on the job, with formal programmes contributing a comparatively modest share. The question pursued here follows from that observation, namely whether reps who have a job aid within reach end up performing better than they otherwise would. In a network

OPEN ACCESS © CINEFORUM

where trainers travel from outlet to outlet, coaching is bound to be uneven. The Job Demands-Resources (JD-R) model (Demerouti et al., 2001) offers one way of framing the problem. Under this model, performance is shaped by the balance between what a job asks of a person and the resources available for meeting it. Read in those terms, a rep at an outlet that trainers visit only now and then is carrying a heavier version of the same job as a colleague elsewhere. AI-driven training would be counted among the resources, since it draws on no trainer time and the rep can open it at any hour. JD-R also implies that the value of such a resource accrues over a period, a reading consistent with Verbeke et al. (2011), whose meta-analysis places work engagement among the drivers of sales performance. Because demands and resources were not measured in this study, the model is applied as an interpretive lens on the observed pattern and not as a hypothesis to be tested, a point taken up again in the Discussion.

3. LITERATURE REVIEW

Twelve studies were reviewed for this paper, and taken together they can be read as three streams of work. Studies in the first stream examine AI in workplace learning and HR. A recurring observation across them is that organisations have adopted AI faster than they have evaluated it, with the benefits claimed for employee development often running larger than the evidence behind them (Bhatt & Muduli, 2023; Pereira et al., 2023; Soulami et al., 2024). Huang and Rust (2021) make a related point in arguing that a more sophisticated AI system will not necessarily perform better. The second stream concerns blended learning in which AI plays no part. Liu et al. (2016) reported in their meta-analysis that blended formats produced better outcomes than no intervention; later research applied the idea to workplace training with self-determination theory as its framework (Kumar et al., 2021; Yu et al., 2025); and Torre et al. (2025), working with a professional training case, observed that learners' satisfaction did not consistently match their learning outcomes. The third stream, covering AI-driven and sales-specific training, is the thinnest of the three and also the closest to the present question, yet it does not settle it. Kim (2022) based the analysis on trainee survey ratings instead of sales figures. Bucher et al. (2024) described a small qualitative pilot. Natto (2026) recorded significantly higher satisfaction under an AI-integrated curriculum, although the gain in academic performance was not significant. Uddin et al. (2026) linked blended learning to performance only through a soft-skills mediator.

Three gaps emerge from this body of work. The first is that most studies of AI-driven or sales-specific training depend on self-reported feedback, whereas the present study works with actual transactional outcomes. The second is that the usual approach compares results before and after an intervention without telling an abrupt shift apart from a longer-run trend; the method adopted here is built to draw exactly that line. The third is that studies tend to report one overall effect and seldom ask whether it holds equally across different kinds of units. This study checks directly whether the pattern is the same across outlets of different sizes.

4. RESEARCH QUESTIONS

These gaps lead to two research questions.

RQ1: Once the AI training tool was in place, did outlet sales performance, measured through net deliveries and tokens, change? If it did, was the change immediate, or did it build up over time?

RQ2: Did any such change hold steady across the three outlets, or did it vary from one outlet to the next?

5. OBJECTIVES

Four objectives turn these questions into specific aims.

To examine whether outlet-level sales performance, in terms of net deliveries and tokens, changed alongside the introduction of the AI training component, using actual transactional records in place of self-reported ratings.

To establish whether any such change took the form of a sudden shift at launch or a gradual build-up over time.

To assess whether the strength of this association was the same across outlets or varied between them.

To add an empirically grounded data point from India to a literature that has so far concentrated outside Indian retail and has seldom used transactional outcome measures.

6. HYPOTHESIS

Two hypotheses carry the first three objectives into testing.

H1: The introduction of the AI-augmented training component is associated with a change in outlet-level sales performance (net deliveries and tokens) over the study period.

H2: The magnitude of this association differs across outlets.

OPEN ACCESS © CINEFORUM

7. CONCEPTUAL FRAMEWORK

Figure 1 sets out these two hypotheses.

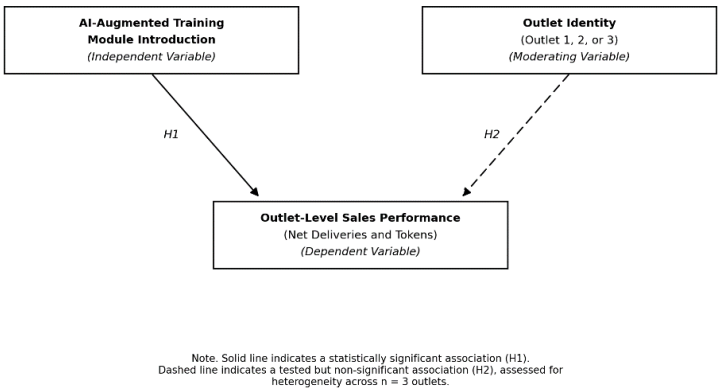


Figure 1. Conceptual framework linking the AI-augmented training module to outlet-level sales performance.
Note. Solid line indicates a statistically significant association (H1). Dashed line indicates a tested but non-significant association (H2), assessed for heterogeneity across n = 3 outlets.
Figure 1 lays out the relationships under test. AI-augmented training, introduced at all three outlets in January 2025, is the independent variable, and outlet-level sales performance is the dependent variable; the solid line joining them represents H1. Outlet identity appears as a moderator on the dashed line and raises the question of whether the size of the association varies by outlet. That interaction turned out not to be significant (Section 10.6). The framework is meant to be read through the Job Demands-Resources lens described in Section 2, in which the AI module counts as an added resource and sales performance is the outcome it is expected to support. JD-R is not given a box of its own, since what the study measures is the transactional pattern alone.

8. OPERATIONAL DEFINITIONS

AI-augmented training component: An on-demand conversational layer, added in January 2025 to the structured six-stage sales training protocol, in which a rep conducts a simulated customer conversation that is then scored on defined competencies.
Net deliveries: The count of gross deliveries once returns have been subtracted. Only sales that the customer did not reverse within the return window are included, so of all the measures in the sales funnel this one is the most conservative, and it is treated as the primary outcome.
Tokens: A reservation made by a customer following a visit to an outlet. Since tokens come earlier in the funnel than deliveries, they serve as the secondary, leading measure, and a training effect, if there is one, would be expected to surface here before it appears in deliveries.
Return rate: Reversed deliveries as a share of gross deliveries. It is a check on the deliveries figure. Should deliveries and returns climb together, a positive reading of the deliveries result would be hard to defend.
Engagement: The count of AI-assessment attempts completed by a sales rep.
Competency score: A rep's highest score across all of their AI-assessment attempts, on the dimensions defined by the platform.
Pre-AI / During-AI period: January–December 2024 and January–December 2025 respectively, on either side of the module's January 2025 launch.

9. METHODOLOGY

9.1 Research Design

Section 9.7 describes the main method, a single-group segmented interrupted time series (ITS) design. Random assignment was off the table, since every outlet adopted the module on the same date. ITS gets around the missing control group by taking the network's own performance before the module as its baseline (Lopez Bernal et al., 2017). Just as important, it can tell a jump at launch apart from a slower bend in the trend. H1 turns on that difference. So does the JD-R argument in Section 2, which expects an added resource to build value gradually and not to lift a demand overnight.

OPEN ACCESS © CINEFORUM

Two supplementary analyses are run next to the ITS model, and they are not interchangeable. The two-way ANOVA handles H2. Its Outlet × Period interaction asks whether the size of the change varies by outlet, something the ITS model has no way of answering. Welch's t-tests with Cohen's d (Section 10.4) are there only to describe within-outlet change in plain terms, beside the ITS coefficients. No hypothesis is tested with them. They are uncorrected for multiple comparisons, and they add no evidence that the ITS and ANOVA results do not already supply. The only test of H2 is the ANOVA in Section 10.6.

9.2 Participants and Setting

In brief, the sample covers N = 3 outlets and 72 outlet-months (36 Pre-AI, 36 During-AI), with 36 sales reps. Only Outlets 1, 2 and 3 in Delhi NCR had data for every month of both periods. Two other outlets were dropped; both opened midway through the During-AI period, after a nearby outlet shut.

9.3 Sampling and Group Assignment

The outlets were chosen by purposive, criterion-based sampling. A random draw made little sense in a network with so few outlets, while convenience sampling would not have ensured complete data across both periods. An inclusion criterion was therefore laid down in advance: an outlet was eligible only where monthly records existed for the whole of the Pre-AI period (January to December 2024) and the whole of the During-AI period (January to December 2025). Outlets that closed, relocated or opened at any point within those two years were excluded under this rule, and three outlets met it.

A sample of three outlets is admittedly small, and it limits the confidence with which differences between outlets can be detected. The non-significant H2 interaction is accordingly interpreted in this paper as a modestly powered result and not as firm proof of uniform impact.

Reps were not sampled separately. Anyone recorded at one of the three outlets is in the data, so each outlet's staff is covered in full, linked through the outlet field in the organisation's employee records (Section 9.4 gives the exact numbers).

9.4 Instruments

Neither data source was built for this study. Both are existing organisational systems, and they are the only ones that routinely generate what was needed: transactional sales figures on one side, AI-assessment records on the other. Data Source 1 is the organisation's standard sales reporting, pulled together in Power BI dashboards. From it came net deliveries, tokens and return rates, based on completed transactions, for all twenty-four months at the three outlets. Records were checked for duplicates, and the data were put on a common calendar so each outlet-month matches up across the two periods. Data Source 2 was the AI-training platform, which logged each rep's attempts, completions and best scores. These were cross-checked against an independent HR roster of 100 employees covering the wider Delhi NCR workforce, of whom 36 work at the three study outlets (8, 17 and 11 at Outlets 1–3 respectively). Matching was done by email address, because the platform's employee-code field does not correspond to the one used in the HR roster.

9.5 Intervention Procedure

The structured six-stage sales training protocol was already in place before the module arrived, delivered in person by trainers who moved from one outlet to another. In January 2025 an on-demand layer was added to it. Reps could now hold a simulated conversation with an AI-led customer, have it scored against fixed competencies, and go through it as many times as they liked. The layer ran in parallel with the in-person sessions and did not take their place.

9.6 Variables

Six variables are used in the study. The independent variable records whether the AI module had launched, that is, whether a given month falls before or after January 2025, and the ITS design is organised around this event. The primary outcomes are net deliveries and tokens, both taken from transactional records instead of self-reported ratings. Tokens register a customer's interest early; net deliveries register the sale once it is done. Both bear on H1. Return rate is a secondary outcome, included to check whether any increase in deliveries came with a matching increase in cancellations.

For H2, outlet identity acts as the grouping variable, since that hypothesis asks whether the size of the association varied between outlets. Engagement and competency are drawn from the 36 relevant AI-platform records. Section 10.6 reports them descriptively to show how reps actually used the tool, and neither is tied to a hypothesis.

OPEN ACCESS © CINEFORUM

9.7 Data Analysis

The main analysis uses the standard segmented ITS regression equation (Lopez Bernal et al., 2017):

$Yt = \beta0 + \beta1(Timet) + \beta2(AI\ Modulet) + \beta3(Time\ Since\ AI\ Modulet) + et$

Here Yt is the outcome in month t, Time is a running month counter, AI Module indicates whether the module had launched (0/1), and Time Since AI Module counts the months elapsed since launch. β2 picks up any sudden jump at launch, while β3 picks up a change in trend, and it is β3 on which this study's finding rests: it was significant for both outcomes, and β2 was not. Newey-West standard errors (two lags) were used to correct for autocorrelation from one month to the next. The segmented ITS regression was fitted to the network-level monthly series (the three outlets' totals added together, over 24 months), whereas the two-way ANOVA and the Welch comparisons drew on the full outlet-month panel.

Both supplementary methods were applied as set out above. The two-way ANOVA tested H2 through its Outlet × Period interaction on the full panel, and Welch's t-tests with Cohen's d supplied the descriptive within-outlet summary. All analyses were run in Python (pandas, statsmodels, scipy.stats).

Cohen's d was calculated as the standardised mean difference between During-AI and Pre-AI values, $d = (M_{(During)} - M_{(Pre)}) / SD_{(pooled)}$. Because net deliveries and tokens are measured in different units, this gives an effect size that can be compared across the two.

9.8 Ethical Considerations

Employees completed the training, AI-assisted component included, as part of their normal job requirements. It was neither started for this study nor changed because of it, and nothing about employees' day-to-day work was altered for research purposes.

Both datasets are extracts of records produced in the ordinary course of the organisation's operations. No information that could identify an individual was kept: outlet names, employee names and internal framework names were all removed before analysis. The values reported here were also transformed using parameters provided by the organisation, and the transformation was designed so that the pattern of results and their statistical significance stayed the same.

The lead researcher works in the training function of the organisation being studied while carrying out this doctoral research on it, and this conflict of interest is disclosed here in full. The anonymisation and transformation steps described above were put in place to manage it. The organisation had no part in the analysis or in the conclusions reached.

10. RESULTS

10.1 Descriptive Statistics

Table 1: Descriptive statistics for the Pre-AI and During-AI periods (outlet-months, n = 36 each)

Metric	Pre-AI 2024 (M, SD)	During-AI 2025 (M, SD)	% Change
Net deliveries	92.22 (30.71)	122.81 (35.61)	+33.2%
Tokens	192.50 (47.82)	263.36 (77.56)	+36.8%
Return rate	4.64%	2.87%	−38.2%

Table 1 presents the figures as they stand before any modelling. Across the 36 outlet-months in each period, mean net deliveries went up from 92.22 to 122.81 (+33.2%) and mean tokens from 192.50 to 263.36 (+36.8%). Over the same comparison, the return rate came down from 4.64% to 2.87% (−38.2%).

10.2 Tests of Assumptions

Before any ITS estimates were interpreted, the assumptions underlying the model were checked. Because monthly business figures tend to carry over from one month into the next, Newey-West HAC standard errors (two lags) were applied. The Durbin-Watson statistic came to 1.77 for net deliveries and 1.28 for tokens, and residual normality was acceptable for both outcomes (Jarque-Bera p = .77 and .71, respectively).

OPEN ACCESS © CINEFORUM

10.3 Outlet-Level Baseline

The outlets did not start out at the same scale. In the months before the AI module was introduced, Outlet 1 averaged 123.00 net deliveries a month, against 86.58 for Outlet 2 and 67.08 for Outlet 3. H2 is therefore judged by how large a change each outlet saw, not by its raw During-AI figures.

10.4 Within-Group Change

Table 2: Net Deliveries, Pre-AI vs During-AI, by Outlet

Group	Pre-AI Mean	During-AI Mean	% Change	Welch t	p	Cohen's d
Outlet 1	123.00	135.50	+10.2%	1.24	.229	0.51
Outlet 2	86.58	127.50	+47.3%	3.50	.003	1.43
Outlet 3	67.08	105.42	+57.1%	3.06	.009	1.25
Network overall	**92.22**	**122.81**	**+33.2%**	**3.90**	**<.001**	**0.92**

Table 2 places the Pre-AI and During-AI means next to each other for every outlet and for the network overall. Outlet 1 moved from 123.00 to 135.50 (+10.2%, p = .229), a change that is not significant on its own. Outlet 2 went from 86.58 to 127.50 (+47.3%, p = .003), and Outlet 3 from 67.08 to 105.42 (+57.1%, p = .009). Across the network, deliveries climbed from 92.22 to 122.81 (+33.2%, p < .001). These per-outlet Welch's t comparisons are descriptive only. They do not make up a formal hypothesis test and are not corrected for multiple comparisons across outlets. H2 itself is tested through the ANOVA reported next, in Section 10.6.

10.5 Primary Analysis

A plain before-and-after comparison cannot tell whether performance jumped at launch or whether its trend shifted, and that is precisely what the ITS model is designed to answer. Both outcomes gave the same answer. There was no sudden shift in January 2025, as the level-change term was not significant. The trend, however, did change. Net deliveries grew by 23.47 units a month over and above the existing trend ($p < .001$, $R^2 = .621$), and tokens by 47.12 a month ($p < .001$, $R^2 = .723$). Such a pattern fits a tool that people took up gradually, though a single-group design cannot exclude other factors that shifted during the same period.

```
In [1]: import pandas as pd, numpy as np
   ...: import statsmodels.api as sm
   ...:
   ...: # Segmented ITS regression - Net Deliveries (network aggregate, 3 outlets)
   ...: X = sm.add_constant(net[['t', 'level', 'trend_post']])
   ...: model_nd = sm.OLS(net["ND's"], X).fit(cov_type='HAC', cov_kwds={'maxlags': 2})
   ...: print(model_nd.summary())
```

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                   ND's   R-squared:                       0.621
Model:                            OLS   Adj. R-squared:                  0.564
Method:                 Least Squares   F-statistic:                     29.58
Date:                Wed, 05 Aug 2026   Prob (F-statistic):           1.49e-07
Time:                        17:42:48   Log-Likelihood:                -128.80
No. Observations:                  24   AIC:                             265.6
Df Residuals:                      20   BIC:                             270.3
Df Model:                           3
Covariance Type:                  HAC
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const        304.9394     26.882     11.344      0.000     252.251     357.628
t             -4.3497      3.489     -1.247      0.213     -11.188       2.489
level         -8.6224     42.364     -0.204      0.839     -91.655      74.411
trend_post    23.4720      6.575      3.570      0.000      10.585      36.359
==============================================================================
Omnibus:                        0.377   Durbin-Watson:                   1.775
Prob(Omnibus):                  0.828   Jarque-Bera (JB):                0.527
Skew:                           0.125   Prob(JB):                        0.769
Kurtosis:                       2.319   Cond. No.                         64.1
==============================================================================

Notes:
[1] Standard Errors are heteroscedasticity and autocorrelation robust (HAC) using 2 lags and without small sample correction
```

Exhibit 1. Python output of the segmented ITS regression (Newey-West HAC, 2 lags) for net deliveries.

OPEN ACCESS © CINEFORUM

10.6 Secondary Outcomes

H2 asks whether the change was of similar size at all three outlets. It was tested with a two-way ANOVA on the full outlet-month panel (72 observations), where the term of interest is the Outlet × Period interaction. For neither outcome was the interaction significant: net deliveries $F(2,66) = 1.87$, $p = .162$, partial $\eta^2 = .054$; tokens $F(2,66) = 2.36$, $p = .102$, partial $\eta^2 = .067$. Running the test on the monthly panel does give it more observations than three outlet means would. The catch is that months from one outlet are not independent, and treating them as if they were makes the test look sharper than it really is. With only three outlets forming the groups, the result is modestly powered and does not show that every outlet responded alike (Section 9.3). Outlet 1, for one, gained less than the other two. A descriptive check set each outlet's engagement and competency scores against its gain. Outlet 1 scored highest but gained least. Outlet 2 scored lowest, yet posted one of the largest gains. The resulting r values, between −0.15 and −0.22, point to essentially no relationship, and with three outlets this is description only, not a statistical test.

```
In [2]: from statsmodels.formula.api import ols
   ...: from statsmodels.stats.anova import anova_lm
   ...:
   ...: # Two-way ANOVA - Outlet x Period interaction (H2), full outlet-month panel
   ...: model = ols('ND ~ C(Outlet) * C(Period)', data=panel).fit()
   ...: aov = anova_lm(model, typ=2)
   ...: print(aov)
```

```
                       sum_sq    df          F    PR(>F)
C(Outlet)          22196.027778   2.0  14.018343  0.000008
C(Period)          16836.125000   1.0  21.266379  0.000019
C(Outlet):C(Period)   2963.083333   2.0   1.871394  0.161984
Residual           52250.750000  66.0        NaN       NaN
```

Exhibit 2. Two-way ANOVA output from Python for net deliveries (Outlet × Period interaction, H2).

10.7 Hypotheses Testing Summary

Table 3: Hypotheses Testing Summary

Hypothesis	Test	Result
H1: AI module associated with change in sales performance	Segmented ITS (trend-change term)	Supported (trend, not level)
H2: Shift magnitude differs across outlets	Two-way ANOVA, Outlet × Period interaction (full monthly panel)	Not supported (n = 3 outlet clusters)

Table 3 puts the two hypotheses side by side. H1 holds, though the support comes from the trend changing and not from any jump at launch. H2 does not hold: on the full monthly panel, the ANOVA found no significant difference in how much the three outlets improved. Three outlet clusters is a thin basis for that test, however, so this is a modestly powered non-finding and not firm proof of uniform impact (Section 9.3).

11. DISCUSSION

Put simply, sales performance at the outlets did not leap when the AI component arrived; what shifted was its trend over the following months. That is roughly what the JD-R reasoning in Section 2 would predict, since an added resource is expected to earn its value gradually and not to clear a demand all at once. Demands and resources were not measured, though, so JD-R remains a way of interpreting the result and has not been tested as a mechanism.

Much of this agrees with the studies reviewed in Section 3. Kim (2022) found that sales training lost effectiveness when the human element was taken out and regained it under a hybrid approach, and this study follows the same

OPEN ACCESS © CINEFORUM

logic by keeping the human-delivered protocol running throughout. Natto (2026) provides a partial contrast, reporting significantly higher satisfaction under an AI-integrated curriculum alongside a difference in academic performance that did not reach significance. The present study, by comparison, followed actual transactional outcomes.

The same caution applies to H2. The 72 outlet-months give the ANOVA more observations than three outlet averages would, yet they are repeated months from the same three outlets and not independent replications. The non-significant interaction therefore cannot establish that all outlets improved to the same degree. Outlet 1 visibly gained less than the other two (Table 2), and confirming this would require a larger sample spanning more organisations.

12. IMPLICATIONS

On the theoretical side, the study adds one of the few evaluations of AI-augmented training identified in Section 3 that rests on transactional outcomes, which speaks directly to the data-modality gap left by self-report studies. It also shows that a segmented ITS design can pull apart an immediate level shift and a gradual change in trend, a distinction that a simple pre-post comparison would have missed.

In practical terms, the benefit seems to accumulate over months of use. Organisations planning similar deployments would do well to set their evaluation timelines with this in mind and hold off on judging early results. The H2 finding, a non-significant and modestly powered difference in pace across outlets, lends tentative support to the view that the benefit is not limited to one type of site, although confirming this needs testing across more outlets.

13. CONCLUSION

This study set out to see whether an AI-augmented training component was linked to a change in outlet-level sales performance at three Delhi NCR outlets, using a segmented ITS design that distinguishes an immediate shift from a slower change in trend. The evidence supports such a link, and for both outcomes it took the form of a sustained change in trend, not a sudden jump. Whether the link differed by outlet was tested directly, and no significant difference was found.

The study contributes in three ways. It measures a real, transactional outcome in place of a self-reported one. It separates a sudden level shift from a slower change in trend. And it does both in an Indian retail setting that this literature has largely left out. Taken as a whole, the work adds a single data point to a body of research that is still small, geographically narrow and mostly built on self-report, and it should be seen as an early phase of a larger effort, not as a final answer

REFERENCES

1. Bhatt, P., & Muduli, A. (2023). Artificial intelligence in learning and development: A systematic literature review. European Journal of Training and Development, 47(7/8), 677–694. https://doi.org/10.1108/EJTD-09-2021-0143
2. Bucher, A., Schenk, B., Dolata, M., & Schwabe, G. (2024). When generative AI meets workplace learning:
3. Creating a realistic & motivating learning experience with a generative PCA. In Proceedings of the European Conference on Information Systems (ECIS 2024). https://doi.org/10.48550/arXiv.2405.15561
4. Demerouti, E., Bakker, A. B., Nachreiner, F., & Schaufeli, W. B. (2001). The job demands-resources model of burnout. Journal of Applied Psychology, 86(3), 499–512. https://doi.org/10.1037/0021-9010.86.3.499
5. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. Journal of the Academy of Marketing Science, 49(1), 30–50. https://doi.org/10.1007/s11747-020-00749-9
6. Kim, S. (2022). Innovating workplace learning: Training methodology analysis based on content, instructional design, programmed learning, and recommendation framework. Frontiers in Psychology, 13, Article 870574. https://doi.org/10.3389/fpsyg.2022.870574
7. Kumar, A., Krishnamurthi, R., Bhatia, S., Kaushik, K., Ahuja, N. J., Nayyar, A., & Masud, M. (2021).
8. Blended learning tools and practices: A comprehensive analysis. IEEE Access, 9, 85151–85197.
9. https://doi.org/10.1109/ACCESS.2021.3085844
10. Liu, Q., Peng, W., Zhang, F., Hu, R., Li, Y., & Yan, W. (2016). The effectiveness of blended learning in health professions: Systematic review and meta-analysis. Journal of Medical Internet Research, 18(1), Article e2. https://doi.org/10.2196/jmir.4807
11. Lopez Bernal, J., Cummins, S., & Gasparrini, A. (2017). Interrupted time series regression for the evaluation of public health interventions: A tutorial. International Journal of Epidemiology, 46(1), 348–355.
12. https://doi.org/10.1093/ije/dyw098

OPEN ACCESS © CINEFORUM

13. Natto, Z. S. (2026). Blended peer-led research curriculum with AI integration improves postgraduate students' academic performance and satisfaction: A quasi-experimental mixed-methods study. BMC Medical Education, 26, Article 260. https://doi.org/10.1186/s12909-026-08576-2
14. Pereira, V., Hadjielias, E., Christofi, M., & Vrontis, D. (2023). A systematic literature review on the impact of artificial intelligence on workplace outcomes: A multi-process perspective. Human Resource Management Review, 33(1), Article 100857. https://doi.org/10.1016/j.hrmr.2021.100857
15. Soulami, M., Benchekroun, S., & Galiulina, A. (2024). Exploring how AI adoption in the workplace affects employees: A bibliometric and systematic review. Frontiers in Artificial Intelligence, 7, Article 1473872.
16. https://doi.org/10.3389/frai.2024.1473872
17. Torre, S., Ulloa Severino, A., & Ligorio, M. B. (2025). Learning outcomes and training satisfaction: A case study of blended customization in professional training. Technology, Knowledge and Learning, 30, 1663–1702. https://doi.org/10.1007/s10758-024-09778-7
18. Uddin, M., Ahamed, M. R., Jakowan, Islam, M. R., & Nahar, M. (2026). How blended learning influences employee performance: The mediating role of soft skill development and knowledge acquisition. International Journal of Accounting and Economics Studies, 13(2), 243–250.
19. https://doi.org/10.14419/9cjynm83
20. Verbeke, W., Dietz, B., & Verwaal, E. (2011). Drivers of sales performance: A contemporary meta-analysis. Have salespeople become knowledge brokers? Journal of the Academy of Marketing Science, 39(3), 407–428. https://doi.org/10.1007/s11747-010-0211-8
21. Yu, Q., Yu, K., & Wang, J. (2025). Unraveling the impact of blended learning vs. online learning on learners' performance: Perspective of self-determination theory. Behavioral Sciences, 15(9), Article 1263.
22. https://doi.org/10.3390/bs15091263

OPEN ACCESS © CINEFORUM