

The Role of Artificial Intelligence in Customer Churn Prediction and Lifecycle Management

¹ Konstantin Zhuchkov

¹ Expert in AI-driven business process optimization New York, USA

Received: 29th Oct 2025 | Received Revised Version: 18th Nov 2025 | Accepted: 29th Nov 2025 | Published: 08th Dec 2025

Volume 07 Issue 12 2025 | Crossref DOI: 10.37547/tajiir/Volume07Issue12-03

Abstract

The article examines how artificial intelligence reshapes churn prediction and customer lifetime value (CLV) management in data-driven marketing. Relevance follows from the growing pressure to align acquisition and retention spending with incremental value under volatile demand and privacy constraints. Novelty lies in integrating uplift-based policy learning, survival-time horizon modeling, and early-signal CLV for media bidding into a single operating blueprint for marketing and sales leaders. The study describes current model families for churn and CLV, reviews evidence on explainability pipelines for managerial sign-off, and evaluates incrementality-aware targeting that reallocates spend from probability-only ranking to causal response. Special attention is given to cold-start acquisition, where CLV must inform bidding before behavioral depth accumulates. The aim is to synthesize actionable guidance that links model choice to levers in pricing, service, offer design, and paid media. Methods include comparative reading, structured content analysis, and decision-matrix synthesis. The conclusion outlines a practitioner-ready cadence for screening, selection, bidding, and governance. The article will benefit CMOs, heads of sales, growth teams, and analytics leaders building durable, value-aligned programs.

Keywords: churn prediction, customer lifetime value, uplift modeling, survival analysis, explainable AI, generative AI, B2B SaaS, media bidding, CAC payback, retention strategy.

© 2025 Konstantin Zhuchkov. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). The authors retain copyright and allow others to share, adapt, or redistribute the work with proper attribution.

Cite This Article: Konstantin Zhuchkov. (2025). The Role of Artificial Intelligence in Customer Churn Prediction and Lifecycle Management. *The American Journal of Interdisciplinary Innovations and Research*, 7(12), 16–22. <https://doi.org/10.37547/tajiir/Volume07Issue12-03>

1. Introduction

Digital customer portfolios expose wide variance in tenure, price sensitivity, and service experience, which undermines static rules for renewal outreach and acquisition pacing. AI-enabled churn screening and individual-level CLV forecasting address this variance by enhancing recall in risk detection, improving horizon estimates for revenue planning, and allocating budgets toward treatments with a positive incremental effect. The purpose of the article is to consolidate recent technical advances into a managerial blueprint that ties model choices to commercial decisions across the lifecycle.

1.1 The study sets three tasks:

- 1 Compare the predictive contribution and managerial transparency of ensemble

classifiers, time-to-event models, and deep sequence architectures for churn and CLV.

- 2 Evaluate the economic impact of incrementality-aware selection in comparison to probability-only ranking and establish guardrails for offline policy evaluation.
- 3 Synthesize an operating cadence that connects early-signal CLV to acquisition bidding and survival-aligned timing to retention, with explainability supporting approval workflows.

Novelty emerges from merging these streams—uplift targeting, survival horizons, and early CLV for paid media—into a single, testable cadence that can be adopted without bespoke infrastructure.

2. Materials and Methods

The synthesis draws on peer-reviewed articles and systematic reviews published between 2021 and 2025. For CLV and financial alignment: N. Ali, O. S. Shabneh [1]. For interpretable, high-recall churn screening: D. Asif, H. Qian, A. Khadidos, A. Rehman, A. Hussain [2]. For incrementality methods and variance-reduction evaluation: B. Bokelmann, A. Bauer, B. Bischl [3]. For B2B SaaS CLV with two-stage temporal modeling: S. Curiskis, X. Dong, F. Jiang, M. Scarr [4]. For a systematic review of CLV algorithms and deployment patterns: E. B. Firmansyah, M. Elveny, A. Affandi, R. Sihombing [5]. For a systematic review of churn models, features, and class-imbalance remedies: M. Imani, V. Kharchenko, M. Aibin [6]. For generative-AI adoption in marketing discovery and creative iteration: N. Kshetri, S. da Cruz [7]. For explainability in telecom churn classifiers: S. S. Poudel, R. Ghosh, D. Basak, S. Das [8]. For benchmarking uplift learners and robust policy estimators: M. Rößler, M. Lachmann, A. Thomas [9]. For multi-output deep CLV in multi-category e-commerce with XAI overlays: G. Y. Benk, B. Badur, S. Mardikyan [10].

Methods: comparative analysis of findings across sources; structured content analysis to identify stable drivers and deployment constraints; synthesis of taxonomies and decision matrices linking models to marketing levers; critical appraisal of offline policy-evaluation practices; integrative reasoning to outline an operating cadence for acquisition and retention.

3. Results

Across recent peer-reviewed work, two consistent outcomes recur: AI models detect churn risk earlier and with greater stability than legacy scoring, and organizations convert these gains into higher lifetime value through targeted retention, pricing, and acquisition bidding. A comprehensive 2025 review organizes modern churn modeling into three families—classical classifiers/ensembles, time-to-event survival models, and deep sequence architectures—and explains how feature construction (usage intensity, contract structure, service issues) and imbalance handling shape real performance.

Figure 1 synthesizes that taxonomy and anchors how model class choices map to managerial levers in retention programs [6]. In parallel, systematic evidence on AI for CLV shows a shift from cohort-level heuristics to individual-level forecasts and prescriptive actions

(e.g., uplift-guided offers and LTV-based targeting) that reallocate budget toward customers with positive, causal response while suppressing spend where incremental value is negative [5].




Fig. 1. Taxonomy of churn-prediction approaches across classical ensembles, survival analysis, and deep sequence models, with typical features and imbalance remedies [6]

Business-facing studies in software and subscription settings suggest that reframing CLV as a supervised learning target—rather than a purely parametric projection—enables richer feature sets and hierarchical ensembling to address heterogeneous segments and data drift. In a B2B SaaS application, a temporal two-stage CLV framework prioritized recent behavior and then corrected it to the whole horizon. Compared to a time-series baseline, the machine-learning framework achieved multi-fold error reductions while preserving managerial interpretability for allocation decisions between acquisition and retention [4]. E-commerce research extended this idea with a multi-output deep network that predicts several value components jointly and exposes feature influences via explainable AI; the result is tighter value segmentation for bid modifiers, free-trial conversion design, and cross-sell sequencing [10]. A 2024 systematic review consolidates these developments, mapping algorithms to data regimes (contractual vs. non-contractual settings; sparse vs. dense telemetry) and reporting that tree-based ensembles and gradient boosting remain strong baselines, while deep models dominate where longitudinal signals are abundant [5].

Explainability requirements—especially in regulated and senior-stakeholder environments—have shifted from “nice-to-have” to a gating factor. Telecom- and platform-data studies show that SHAP-style explanations recover

stable, business-plausible drivers of churn (tenure discontinuities, month-to-month contracts, price jumps, service-issue proxies), improving trust and accelerating remediation design without sacrificing accuracy [8]. Large-scale work coupling gradient-boosted ensembles with transparent rule extraction or post-hoc explanations demonstrates that high-recall churn screening can co-exist with auditable reasoning and workload triage for care teams [2]. These interpretability pipelines reduce the back-and-forth between analytics and marketing/sales by making the “why” of risk immediately legible for creative and offer design [6; 8].

A second, decisive pattern is the rise of incrementality-aware decisioning. Traditional retention targeting based on predicted churn probability alone often wastes offers on customers who would have stayed or leaves value on the table among those who would only stay with a specific treatment. Recent evaluations in operations research compare uplift-modeling strategies (two-model, meta-learner, and direct-uplift learners) under variance-reduction estimators; the studies find materially better policy value when selection is driven by incremental response rather than risk alone, with clear guidance on when each learner is robust (sample size, noise level, overlap) [3]. A separate benchmarking stream emphasizes the need for consistent, counterfactual-aware metrics and demonstrates that policy-evaluation choices can alter model rankings; adopting these robust estimators stabilizes model selection across industries and campaign cadences [9]. For a marketing & sales organization, these results translate into concrete playbooks: gate all retention spend through uplift-rankers, use robust off-policy evaluation during design weeks, and reserve expensive treatments for segments with both high uplift and high projected contribution margin [3; 9].

On the acquisition side—where LTV must be forecast before much behavior accumulates—recent work integrates early-signal CLV predictions into bidding and audience construction. The literature reports operational gains from feeding model-based CLV into media platforms as a conversion proxy or as a post-bid adjustment, which aligns customer acquisition cost (CAC) to downstream value and shortens payback windows [4; 5]. Studies in innovation and marketing science document how generative-AI-enabled market research compresses qualitative discovery cycles and increases campaign productivity, allowing teams to iterate on audience hypotheses, creative variants, and

offer frames at a fraction of the historic timelines while preserving human validation for brand and compliance guardrails [7]. For organizations that pair this accelerated research with uplift-screened experiments, the compounded effect is a tighter test-learn loop: faster discovery of segment pain points, quicker creative fit, and more efficient reallocation between acquisition and retention pools [3; 7].

Time-to-event modeling fills a specific lifecycle gap: when the timing of churn carries economic weight (renewal cycles, seasonal offboarding), survival models provide hazard curves that inform staffing and inventory planning, while improving CLV projections through better lifetime horizons [6; 8]. Evidence reviews show survival approaches outperform static classifiers on datasets with censoring and staggered starts, and that hybrid pipelines—survival for horizon, gradient boosting for treatment effect, and policy optimization for offer choice—deliver the most stable financial outcomes across quarters [6; 8; 9]. In data-rich e-commerce, multi-task networks generalize this idea by predicting near-term value, long-term value, and cross-sell propensity together, which improves bidding and sequencing logic at the funnel top and mid-funnel [10].

Taken together, these findings outline a practical operating model for marketing and sales leaders. First, standardize an explainable, class-imbalance-aware churn stack for high-recall screening that hands off intelligible drivers to service and retention design [2; 6; 8]. Second, replace probability-only selection with uplift-based policies and robust offline evaluation, ensuring offer spend purchases incremental LTV rather than raw risk reduction [3; 9]. Third, propagate early CLV predictions into acquisition bidding and look-alike construction, so that CAC dynamically reflects the expected contribution. This is achieved by utilizing generative AI-accelerated market research to keep hypotheses and creative fresh at a low marginal cost [5; 7]. Finally, use time-to-event outputs to refine lifetime horizons and staffing plans, thereby closing the loop by monitoring realized margin and incorporating feedback into both churn and CLV models.

4. Discussion

AI-driven churn screening and individual-level CLV forecasting alter marketing and sales planning by shifting the selection from heuristic rules to model-based policies that prioritize incremental value over sheer risk

reduction. Across reviewed studies, gradient-boosted ensembles and other tree methods remain dependable for tabular, mixed-quality CRM data. Survival analysis improves horizon estimates where renewal timing matters, and deep sequence architectures add lift when longitudinal clickstream or telemetry data are dense [4–6; 8; 10]. At the commercial layer, CLV-informed bidding aligns acquisition with downstream contribution, reducing CAC payback volatility. In subscription and B2B SaaS settings, early CLV models inform audience construction and budget pacing, resulting in measurable error reductions against time-series baselines [1; 4; 5].

Generative AI tooling shortens audience research cycles and creative iteration, provided that outputs are validated with uplift-aware experiments rather than proxy engagement metrics [3; 7; 9].

To translate model fit into reliable action, the modeling choice must be tied to the marketing lever it unlocks. Table 1 consolidates the decision surface reported in the literature and links each approach to the specific operational move it enables in retention and acquisition programs.

Table 1 – Mapping of churn/CLV modeling choices to marketing and sales decision levers (Compiled by the author based on sources: [1; 3–10])

Modeling approach	Data regime fit	Primary decision lever	Implementation note
Gradient-boosted ensembles / tree-based classifiers	Tabular CRM with mixed sparsity; moderate history length	High-recall churn screening for service triage and offer gating	Calibrate with class-imbalance remedies and post-hoc explanations for managerial sign-off
Time-to-event (survival/hazard) models	Contractual settings with censoring and renewal cycles	Timing of interventions and lifetime horizon refinement feeding CLV	Couple hazard outputs with cost curves to schedule outreach and staffing
Deep sequence / multi-output networks	Dense behavioral telemetry; multi-category e-commerce	Joint prediction of near-term value, long-term value, and cross-sell sequencing	Use XAI overlays to expose drivers for creative and merchandising teams
Uplift learners (two-model, meta-learners, direct)	Campaign data with treatment/control and sufficient overlap	Selection on incremental response, not risk alone, for retention spend	Evaluate with variance-reduction estimators; reserve high-cost offers for high-uplift/high-margin segments
Early-signal CLV for acquisition bidding	Cold-start or short histories in paid media	CAC bids proportional to expected contribution; audience pruning	Pipe model outputs as conversion proxies or post-bid adjustments; monitor payback windows
GenAI-supported audience research	Rapid hypothesis generation for segments and creative	Faster test-learn cycles that feed uplift experiments	Keep human review and brand/compliance guardrails; measure by incremental value

Post-deployment experience in telecom, platforms, and retail underlines that interpretability is not a cosmetic add-on. SHAP-style explanation pipelines consistently surface business-plausible churn drivers (such as contract type, tenure, price jumps, and service incidents), thereby decreasing the friction between analytics and channel owners and accelerating the redesign of offers and journeys without measurable accuracy loss [2; 6; 8]. Survival outputs further reduce planning noise by providing hazard-aligned windows for outreach, which in turn stabilizes CLV projections and resource allocation to retention pods [8]. For acquisition, the combination of early-signal CLV and uplift-tested creatives narrows audience uncertainty and supports

disciplined budget pacing at weekly and intra-week horizons [4; 5; 7; 9].

Evidence on incrementality reshapes the standards for target selection. Probability-only ranking tends to overspend on “self-retainers” and miss segments that respond only under specific treatments. Comparative evaluations demonstrate that uplift learners, when evaluated with robust policy estimators, increase policy value and prevent reversals in offline selection that result from naïve metrics [3; 9]. Table 2 groups the recurrent implementation risks that degrade financial impact and notes the mitigations that the literature associates with durable performance.

Table 2 – Recurrent implementation risks in churn/CLV programs and reported mitigations (Compiled by the author based on sources: [1-9])

Challenge	Observed effect in studies	Mitigation reported
Data shift and drift across quarters	Decay in recall/precision; unstable ROI	Rolling refits; drift monitoring; champion-challenger with cost-sensitive thresholds
Severe class imbalance	Over-confident negatives; missed at-risk cohorts	Stratified sampling, calibrated probabilities, cost-weighting; recall-first thresholds
Label leakage from post-treatment fields	Inflated offline AUC; poor generalization	Leakage audits; feature governance; time-safe feature generation
Limited stakeholder trust in black-box models	Slower adoption; blocked campaigns	XAI (e.g., SHAP) with stability checks and rule extraction; documentation for audit
Biased offline policy evaluation	Model ranking inversions; wasted spend	Variance-reduction and doubly robust estimators; overlap diagnostics
Misaligned CAC in paid media	Long payback; budget whiplash	Early-signal CLV as proxy objective or post-bid multiplier; weekly payback reviews
Creative/testing throughput bottlenecks	Stagnant uplift; fatigued teams	GenAI-assisted ideation with human QA; fast A/B gating by incremental value

For a marketing and sales organization, these findings support a concrete operating cadence. Churn screening

should run on calibrated ensembles with explanation dashboards that expose stable drivers for service and

pricing squads, while survival outputs inform renewal-adjacent staffing and cadence [2; 6; 8]. Retention spend selection is attached to uplift rankers evaluated using doubly robust estimators; policy rollouts proceed in waves sized by confidence intervals rather than point estimates [3; 9]. Acquisition teams prioritize early-signal CLV as the optimization target or post-bid multiplier and prune audiences that exhibit negative expected contribution, which ties CAC to contribution in real time [1; 4; 5]. GenAI shortens audience research cycles and increases creative diversity; however, campaign promotion rights remain contingent on incremental value observed in tests, rather than on proxy engagement [7; 3; 9].

Limitations arise from privacy constraints, sparse treatment data in early programs, and the operational costs associated with continuous evaluation. Studies emphasize that generalization suffers when overlap between treated and control cohorts is weak, implying that marketing operations need deliberate exploration to sustain uplift learning rather than one-off pilots [3; 9]. Reviews of CLV modeling caution against cross-domain extrapolation without re-estimating lifetime horizons and contribution margins, particularly when seasonality and contract structures differ [5; 6]. Telecom and platform evidence on explainability indicates that explanations must be audited for stability across resamples; unstable attributions can mislead journey redesigns [2; 8]. These constraints do not negate the reported gains; they define the governance and measurement scaffolding that preserves financial impact at scale while keeping teams focused and workloads sustainable.

5. Conclusion

Comparison of model families—shows that calibrated gradient-boosted ensembles provide dependable high-recall churn screening on tabular CRM, survival models improve renewal-adjacent timing and lifetime horizons, and deep sequence or multi-output networks add lift where longitudinal telemetry is dense; explainability overlays deliver stable, business-plausible drivers for authorization and redesign workflows.

Incrementality versus probability—confirms that uplift-based selection and doubly robust evaluation raise policy value relative to probability-only ranking, reduce wasted offers on self-retainers, and stabilize campaign rollout decisions under finite test budgets.

Operating cadence—yields a practitioner blueprint: run explainable, imbalance-aware churn screening for triage; use uplift rankers to gate retention spend; propagate early-signal CLV into bidding and audience pruning to align CAC with downstream contribution; schedule interventions with survival-aligned windows; institutionalize governance for drift monitoring, leakage audits, and explanation stability. The resulting program links analytics to commercial levers in pricing, service, media, and creative, enabling faster growth with measured workload and auditable financial impact.

References

- 1 Ali, N., & Shabn, O. S. (2024). Customer lifetime value (CLV) insights for strategic marketing success and its impact on organizational financial performance. *Cogent Business & Management*, 11(1), 2361321. <https://doi.org/10.1080/23311975.2024.2361321>
- 2 Asif, D., Arif, M. S., & Mukheimer, A. (2025). A data-driven approach with explainable artificial intelligence for customer churn prediction in the telecommunications industry. *Results in Engineering*, 26, 104629. <https://doi.org/10.1016/j.rineng.2025.104629>
- 3 Bokelmann, B., & Lessmann, S. (2022). Improving uplift model evaluation on RCT data: A preprint (pp. 1–22). arXiv. <https://arxiv.org/pdf/2210.02152>
- 4 Curiskis, S., Dong, X., Jiang, F., & et al. (2023). A novel approach to predicting customer lifetime value in B2B SaaS companies. *Journal of Marketing Analytics*, 11, 587–601. <https://doi.org/10.1057/s41270-023-00234-6>
- 5 Firmansyah, E. B., Machado, M. R., & Moreira, J. L. R. (2024). How can Artificial Intelligence (AI) be used to manage Customer Lifetime Value (CLV)—A systematic literature review. *International Journal of Information Management Data Insights*, 4(2), 100279. <https://doi.org/10.1016/j.jiimei.2024.100279>
- 6 Imani, M., Joudaki, M., Beikmohammadi, A., & Arabnia, H. R. (2025). Customer churn prediction: A systematic review of recent advances, trends, and challenges in machine learning and deep learning. *Machine Learning and Knowledge Extraction*, 7(3), 105. <https://doi.org/10.3390/make7030105>
- 7 Kshetri, N., Dwivedi, Y. K., Davenport, T. H., & Panteli, N. (2024). Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda. *International*

Journal of Information Management, 75, 102716.
<https://doi.org/10.1016/j.ijinfomgt.2023.102716>

8 Poudel, S. S., Pokharel, S., & Timilsina, M. (2024). Explaining customer churn prediction in telecom industry using tabular machine learning models. *Machine Learning with Applications*, 17, 100567. <https://doi.org/10.1016/j.mlwa.2024.100567>

9 Rößler, J., & Schoder, D. (2022). Bridging the gap: A systematic benchmarking of uplift modeling and heterogeneous treatment effects methods. *Journal of Interactive Marketing*, 57(4), 629–650. <https://doi.org/10.1177/10949968221111083>

10 Yilmaz Benk, G., Badur, B., & Mardikyan, S. (2022). A new 360° framework to predict customer lifetime value for multi-category e-commerce companies using a multi-output deep neural network and explainable artificial intelligence. *Information*, 13(8), 373. <https://doi.org/10.3390/info13080373>

Figure




Fig. 1. Taxonomy of churn-prediction approaches across classical ensembles, survival analysis, and deep sequence models, with typical features and imbalance remedies [6]