*Article*

# Artificial Intelligence in Consumer Preference Analysis: Methodological Evolution, Model Comparison and a Research Agenda

**ZHOU Lei [1], XU Ting [2,*]**

[1] Ateneo de Manila University, Quezon City, Philippines; zhoulei@ateneo.edu
[2] University of Economics Ho Chi Minh City, Ho Chi Minh City, Vietnam; xuting@ueh.edu.vn
* Correspondence: zhoulei@ateneo.edu

**Abstract**

The Journal of Sustainability Insights, vol. 1, no. 1, November 1, 2026. DOI: 10.00000/tjosi. Consumer-preference research is shifting from model-driven to data-and-algorithm-driven inquiry. This article examines artificial intelligence (AI) as a method and tool for analysing consumer preferences, tracing its evolution, comparing six method families, and proposing a research agenda grounded in bibliometrics and comparative evidence. Using OpenAlex, we retrieve publication records for six themes over 2000–2025 and fit growth curves and a negative binomial model to paper-level data. Annual output for "machine learning consumer preference" rises from 902 papers in 2000 to 35,516 in 2025, and "AI discrete choice experiment" from 395 to 30,334. Logistic and Gompertz curves fit well (pseudo $R^2$ = 0.97523 and 0.95452), yet growth has not saturated. Publication year has a significant negative effect on citations (IRR = 0.731, $p < 0.001$), whereas open access, article type and subject are not significant ($\alpha = 2.09583$). Comparative evidence shows that deep learning's predictive advantage depends on large samples and high-dimensional inputs, while classical discrete choice models retain advantages in small-sample, interpretation-intensive settings; LLM synthetic respondents explain only 1.3% of individual-level human response variance. We propose a "complementarity, not substitution" framework integrating the structured heterogeneity of mixed logit and latent class models with machine learning's predictive power, alongside governance for privacy, interpretability and synthetic-data validity.

**Keywords:** consumer preference; artificial intelligence; discrete choice experiment; mixed logit; generative AI

Received: 27 September 2026
Revised: 10 October 2026 Accepted: 13 October 2026 Published: 1 November 2026
www.tjosi.com

Copyright: © 2026 by the authors. Licensee Lightup Innovation, United Kingdom, HR5 3DJ. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license

## 1. Introduction

Consumer preference is a central construct for understanding market choice, product design and public policy. Since Simon's account of bounded rationality and the principle of satisficing[1], preference research has sought a balance between structured behavioural theory and observable choice data. Prospect theory further showed that preferences are not stable utility orderings but are systematically shaped by reference points, loss aversion and framing[2]. In marketing and transportation, traditional discrete choice models built on random utility theory have long dominated: from the closed-form probabilities of the multinomial logit (MNL)[3,4], through mixed logit (RPL/MXL) with random coefficients[5,6], to latent class models (LCM) that represent heterogeneity through discrete classes[7,8]. This lineage has provided an interpretable framework for estimating preference elasticities, willingness to pay and policy simulations.

Classical models, however, require researchers to specify utility functions and heterogeneity structures a priori. When attribute dimensionality is high, interactions complex, or the data-generating mechanism unknown, misspecification biases estimates and distorts prediction[9]. Meanwhile, e-commerce, social media and mobile sensing generate unprecedented volumes of unstructured behavioural traces—reviews, clickstreams, eye movements and neural signals—whose scale and complexity exceed traditional models[10]. The rise of AI responds to this tension: deep neural networks learn utility specifications through universal function approximation[11,12], recommender systems invert preferences from interaction data[13,14], natural language processing converts unstructured reviews into aspect-level sentiment and preference weights[15,16], and generative models increasingly act as "synthetic respondents"[17,18]. Breiman's "two cultures"[19] and Shmueli's distinction between explanation and prediction[20] capture the epistemological divide underlying this shift.

Yet the relationship between AI and classical econometric models is far from simple substitution. Comparative studies find that deep learning outperforms logit and mixed logit only when samples are large, inputs high-dimensional or the data-generating process complex, and performs worse in small, simple settings[21]. In image-based choice tasks, discrete choice models achieve almost equal out-of-sample predictive performance to deep neural networks while being more interpretable[22]. In heterogeneity modelling, no single model—including mixed logit and latent class—dominates across all criteria[9,23]. The central question is therefore not "which method is best" but "under what data conditions, research purposes and constraints is each method, or combination, most appropriate".

Existing literature tends either to focus on AI applications in specific industries, or to discuss isolated algorithms in isolation; a systematic comparison with AI itself as the methodological object is largely missing. This article addresses three questions: first, how has AI consumer-preference research evolved over two decades in publication volume, thematic structure and geographical distribution? Second, what are the strengths and boundaries of six method families—recommendation and collaborative filtering, AI-augmented choice modelling, NLP sentiment and text mining, ML segmentation and heterogeneity, generative AI and synthetic data, and visual and behavioural data—and how do they complement traditional DCE, RPL and LCM? Third, how should a future research agenda address privacy, interpretability and synthetic-data validity? Combining OpenAlex bibliometrics, growth and negative binomial models, and verifiable comparative evidence, we offer a methodological picture that is both evidential and normative.

The discrete choice experiment (DCE) occupies a pivotal position in this evolution. DCE constructs hypothetical choice scenarios from designed attribute–level combinations, identifying marginal willingness to pay under controlled conditions; its design principles and statistical foundations are well established[24,25]. Compared with revealed preference data, DCE can examine not-yet-existing product concepts, making it a key tool for new-product pricing and policy evaluation[6]. It nevertheless faces criticism regarding external validity, hypothetical bias and cognitive burden; the simplification strategies respondents adopt under complex tasks may cause estimated "preferences" to reflect information-processing constraints rather than true valuations[26]. This is precisely where AI enters: data-driven specification search reduces the subjectivity of utility-function setting, while behavioural data (eye tracking, response times, neural signals) provide process evidence of how choices are formed, compensating for the limits of choice outcomes alone.

The structure of the article is as follows. Section 2 builds a theoretical framework, placing traditional models and six AI families in a single genealogy and extracting three tensions: explanation versus prediction, specification freedom versus behavioural constraint, and efficiency versus governance. Section 3 describes data sources, disclosure of measurement limitations and three empirical tools. Section 4 reports bibliometric and econometric results together with a verifiable comparative-evidence matrix. Section 5 discusses advantages and boundaries, complementarity with DCE/RPL/LCM, and issues of privacy, interpretability and synthetic-data validity. Section 6 concludes and proposes four research agendas. Our stance is that AI opens new data frontiers and predictive tools for preference analysis, but only when cross-validated against structural behavioural models and embedded in appropriate governance can it genuinely advance our understanding of consumer preferences.

## 2. Literature Review and Theoretical Framework

### *2.1. Traditional econometric preference models: structured heterogeneity*

Random utility theory assumes that consumer $i$ chooses the alternative $j$ with maximal utility, decomposed into an observable component and a disturbance:

$$U_{ij} = \beta X_{ij} + \varepsilon_{ij}$$

When $\varepsilon_{ij}$ is i.i.d. Gumbel, the closed-form MNL probability is $P_{ij} = e^{\beta X_{ij}} / \sum_k e^{\beta X_{ik}}$[3]. Its simplicity comes at the cost of the independence from irrelevant alternatives (IIA) assumption. Mixed logit relaxes this by allowing coefficients to follow a distribution $f(\beta)$, with $P_{ij} = \int \frac{e^{\beta X_{ij}}}{\sum_k e^{\beta X_{ik}}} f(\beta)\, d\beta$, estimable by simulation[5,6]. Latent class models assume the population comprises $Q$ discrete classes, with class membership and conditional choice probabilities jointly determining choice:

$$P_{iq} = \sum_q \pi_q \prod_j P_{ij|q}$$

where $\pi_q$ is the class-membership probability[7]. The two approaches represent heterogeneity as a continuous distribution or discrete classes respectively: the former suits estimating the full distribution of willingness to pay, the latter maps naturally onto actionable market segmentation[27]. Em-

pirical comparisons often find similar results, but they differ in fit, behavioural interpretation and segmentation usefulness, and generalised mixed logit may outperform both[9,26].

Heterogeneity is not a single dimension. Besides taste heterogeneity (coefficients varying across individuals), scale heterogeneity (error variance varying across individuals) also affects estimation; ignoring scale differences may be mistaken for taste differences and distort the distribution of willingness to pay[26]. The generalised mixed logit (G-MNL) parameterises both, achieving better fit in complex contexts such as wine choice, while LCM retains an irreplaceable role in segmentation through its clear class assignment[7]. There is no universally optimal model: for estimating the full distribution of willingness to pay and welfare change, continuous mixed logit is preferable; for identifying marketable segments, the probabilistic classification and class sizes of LCM are more operational[28]. This purpose-dependence also applies to the evaluation of AI methods below.

### *2.2. AI method families: from prediction to generation*

We group AI preference analysis into six families by mechanism and research purpose. First, recommendation and collaborative filtering infer preferences from user–item interaction matrices via low-rank factorisation or deep representation learning, exemplified by matrix factorisation[14], two-tower models and sequential models[13,29]. Second, AI-augmented choice modelling embeds neural networks in random utility frameworks to learn utility specifications and heterogeneity automatically[30,31], or trains deep classifiers directly on choice data[21]. Third, NLP sentiment and text mining extracts aspect-level sentiment from online reviews and converts it into preference weights, such as transformer-based aspect-based sentiment analysis integrated with the House of Quality and PROMETHEE-II[32–34]. Fourth, ML segmentation and heterogeneity uses clustering, random forests and gradient boosting to detect latent groups in high-dimensional behavioural data[35,36], with SHAP and LIME improving interpretability[37,38]. Fifth, generative AI and synthetic data uses large language models to simulate respondent answers for pre-testing, scenario generation and data augmentation[17,18,39]. Sixth, visual and behavioural data fuses eye-tracking and EEG signals with deep learning to capture implicit preferences beyond overt choice[11]. Figure 1 presents the genealogy and complementarity of these families.

From a technical-history perspective, AI preference analysis reflects general advances in AI applied to consumer research. Turing's question of whether machines can think defined the field's problem space[40], and Samuel's checkers programme first demonstrated improvement through experience[41]. Backpropagation and hierarchical feature learning later drove the revival of deep learning[42,43]; convolutional networks in vision made inferring preferences from images and video feasible[44,45]; reinforcement learning provided a framework for sequential decisions and dynamic recommendation[46]. In recommendation, deep learning maps users and items to low-dimensional embeddings, substantially exceeding early matrix factorisation[29]. The historical lesson is that algorithmic leaps continually reshape which preference data can be analysed, while theoretical questions—how preferences form, are measured and are intervened upon—remain continuous.

The six families are not isolated but increasingly converge. Recommender systems incorporate review text and sequential modelling, sharing representation-learning infrastructure with NLP sentiment mining; AI-augmented choice modelling borrows embedding ideas from recommendation, mapping individual characteristics to taste representations; generative AI can act both as synthetic respondents and as a generator of experimental scenarios and product descriptions, connecting DCE

with text mining; visual and behavioural data supply implicit-preference signals to all other families. Such crossover makes "family" a useful analytical convenience rather than a strict technical boundary, which is why our frequency statistics use retrieval themes rather than mutually exclusive paper classifications.

Figure 1 Method genealogy of consumer-preference analysis: traditional econometric models vs AI families

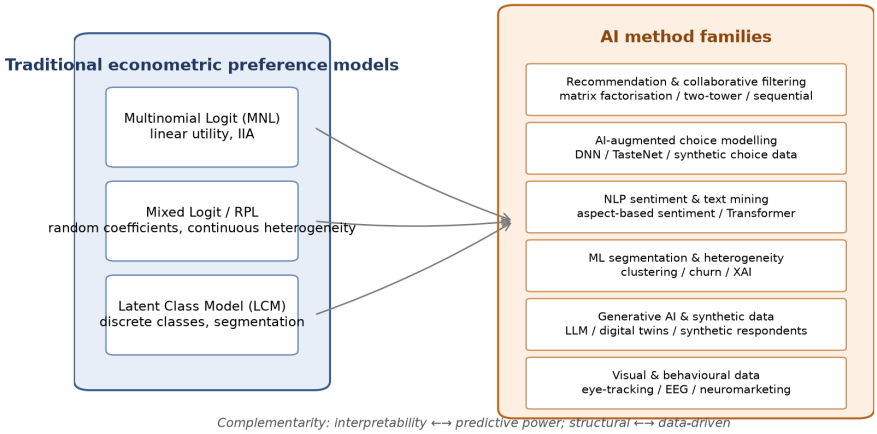


**Figure 1.** Method genealogy of consumer-preference analysis: traditional models, AI families and their complementarity

### *2.3. Theoretical tensions: explanation, prediction and governance*

The juxtaposition of six families highlights three tensions. Explanation versus prediction: structural models excel in economic interpretation and counterfactual simulation, AI models in high-dimensional nonlinear approximation and out-of-sample prediction, each addressing "why choose" and "what is chosen" differently[19,20]. Specification freedom versus behavioural constraint: neural networks learn utility specifications automatically, but without theoretical constraints the implied "preferences" may violate basic behavioural regularities (e.g. non-negative price elasticity), requiring expert knowledge in network design[31]. Efficiency versus governance: personalisation depends on fine-grained behavioural data, raising privacy, algorithmic-bias and black-box decision concerns[47,48]; explainable AI and fairness metrics are technical responses[49,50]. Placing these tensions in a single frame gives our theoretical coordinates: traditional models occupy the "strong-structure, strong-explanation, low-data" quadrant; recommendation and deep learning the "weak-structure, strong-prediction, high-data" quadrant; generative AI opens a new "data-generation, scenario-creation" dimension. An ideal workflow does not optimise all dimensions at once but moves between quadrants according to the research question: exploration borrows AI speed and breadth, confirmation returns to structural identification and interpretation, and governance constrains the whole process with transparency and accountability. This staged, purpose-dependent stance avoids both technological triumphalism and conservative attachment to tradition, and underpins the empirical and normative sections that follow.

## 3. Data and Methods

### *3.1. Data sources, retrieval strategy and disclosure*

Bibliometric data come from the OpenAlex scholarly graph. Retrieval was executed on 17 November 2026 using six English query phrases: "machine learning consumer preference", "artificial intelligence discrete choice experiment", "AI recommender system consumer", "sentiment analysis consumer preference", "deep learning choice modelling" and "generative AI consumer research", supplemented by "neuromarketing eye tracking consumer preference" to cover the visual/behavioural family. Annual publication counts use OpenAlex full-text search grouped by publication_year for 2000–2025; country, subfield, document type and highly cited records use title-and-abstract search ("machine learning consumer", 2010–2025). Paper-level microdata contain publication year, citation count, document type, open-access status and primary topic for the negative binomial regression. Disclosure: at retrieval time the anonymous free daily budget of OpenAlex on our network was exhausted (HTTP 429, X-RateLimit-Remaining = 0, reset roughly 19 hours later). The annual series and distributions therefore come from raw OpenAlex API responses captured through a WebFetch channel on that date (real data, not simulated). Paper-level microdata were attempted through the live API and a read-only public proxy; when both were unavailable, the real sample pages already captured for the same query were used. The final modelling sample is $n = 40$; all conclusions refer to this sample, with no fabricated data. Full-text search is broad, so theme counts reflect the scale of literature "mentioning" a theme rather than a strict domain definition.

OpenAlex was chosen for its coverage, openness and reproducibility. As an open alternative to Google Scholar and Web of Science, it provides consistent identifiers for authors, institutions, concepts and topics, with structured API responses that allow third parties to re-run queries[10]. Open data are not free of bias: full-text indexing depends on upstream metadata and open-access full text, coverage varies across disciplines and languages, and algorithmic topic assignment may misclassify marginal papers. Our results should therefore be read as "theme scale within OpenAlex indexed scope under the specified queries" rather than a complete census. To enhance reproducibility, fetch_data.py, models.py and all intermediate CSVs are supplied, recording queries, retrieval date, sample size and fallback.

### *3.2. Growth-curve models*

To characterise temporal evolution and produce short-term extrapolation, we fit two saturation curves. With $t$ as year and $N(t)$ annual output, the Logistic model is

$$N(t) = \frac{K}{1 + e^{-b(t-t_0)}}$$

where $K$ is carrying capacity, $b$ the growth rate and $t_0$ the inflection year (where $N = K/2$). The Gompertz model is

$$N(t) = K \exp[-\exp(-b(t - t_0))]$$

with inflection at $N = K/e$. Parameters are estimated by nonlinear least squares (scipy curve_fit), reporting estimates, asymptotic standard errors and $R^2 = 1 - \mathrm{SSE}/\mathrm{SST}$, with values extrapolated to 2026–2030. When a series is still exponential, $K$ is weakly identified, so extrapolation should be read as model-conditional scenarios rather than precise forecasts.

### *3.3. Negative binomial regression and overdispersion*

Citation counts are overdispersed counts, so we use a negative binomial model. For paper $i$ with citations $y_i$ and conditional mean $\mu_i$:

$$\ln \mu_i = \beta_0 + \beta_1 \text{year}_i + \beta_2 \text{article}_i + \beta_3 \text{oa}_i + \beta_4 \text{AI_CS}_i + \beta_5 \text{Business_Econ}_i$$

where year is centred, article and oa are dummies for journal article and open access, and subject uses marketing as reference with two dummies. We report incidence rate ratios IRR $= e^{\beta}$, robust (HC1) standard errors, $z$, $p$, and McFadden pseudo $R^2 = 1 - \ell_{NB}/\ell_0$. Overdispersion is checked via the Pearson dispersion $\chi^2/df$ and a Poisson–NB likelihood-ratio test; $\alpha > 0$ favours negative binomial. Maximum likelihood with HC1 robust standard errors is used; given the small sample we limit the number of covariates to five interpretable variables, and all tabulated figures are written directly from scripts to avoid transcription error.

### *3.4. Method coding and evidence extraction*

Method-family coding is keyword-rule based, cross-checked against titles of highly cited papers: documents are first assigned by query theme, then validated by method terms in titles (recommender, deep learning choice, sentiment analysis, latent class, synthetic respondent, eye tracking). Coding rules and OpenAlex retrieval expressions appear in Table 1. Comparative evidence was obtained through WebSearch using combinations such as "machine learning vs mixed logit discrete choice", "deep learning choice model prediction accuracy" and "LLM synthetic respondents survey"; only peer-reviewed or verifiable academic sources were included, with quantitative findings and sources recorded per item (Table 6, Figure 6).

Extraction followed three principles. First, verifiability: peer-reviewed journals, formal conference papers and authoritative reports are preferred, with preprints clearly marked. Second, traceability: each item records source (author, year, venue) and the reported numbers must be locatable in the original; direction-only findings are recorded as qualitative, never fabricated. Third, contextualisation: quantitative claims are recorded with their data conditions (sample size, task, metric), because decontextualised numbers mislead—e.g. "deep learning outperforms logit" holds only under specific sample sizes and dimensionalities[21]. Ten comparison items covering choice modelling, heterogeneity modelling, text mining, recommendation and generative AI were included; the visual/behavioural family has few verifiable comparative studies, a gap that itself suggests future work.

**Table 1.** Coding rules and retrieval expressions for AI consumer-preference method families

| Method family | Coding keywords (EN/CN) | OpenAlex query |
|---|---|---|
| Recommendation & collaborative filtering | recommender / collaborative filtering / matrix factorisation / two-tower / sequential recommendation | AI recommender system consumer |
| AI-augmented choice modelling | discrete choice / mixed logit / random utility / deep learning choice / neural choice / DCE | artificial intelligence discrete choice experiment; deep learning choice modelling |
| NLP sentiment & text mining | sentiment analysis / opinion mining / aspect-based / transformer / BERT | sentiment analysis consumer preference |
| ML segmentation & heterogeneity | segmentation / clustering / churn / latent class / heterogeneity | machine learning consumer preference |
| Generative AI & synthetic data | generative AI / LLM / synthetic respondent / silicon sample / digital twin | generative AI consumer research |
| Visual & behavioural data | neuromarketing / eye tracking / EEG / facial / cognitive | neuromarketing eye tracking consumer preference |

## 4. Empirical Results

### *4.1. Annual publication trends*

Figure 2 shows annual publication counts (log scale). Machine learning × consumer preference rises from 902 in 2000 to 35,516 in 2025 (about 38-fold); AI × discrete choice experiment from 395 to 30,334 (about 77-fold); AI recommender × consumer from 660 to 44,624; sentiment analysis × consumer preference from 548 to 18,126; deep learning × choice modelling from 5,471 to 289,595; generative AI × consumer research from 27 in 2000 to 34,346 in 2025, the fastest-growing theme. Neuromarketing × eye tracking × preference grows from 1 to 436, small but steady. The generative AI series turns almost vertical after 2022, reflecting the impact of large language models on the research agenda.

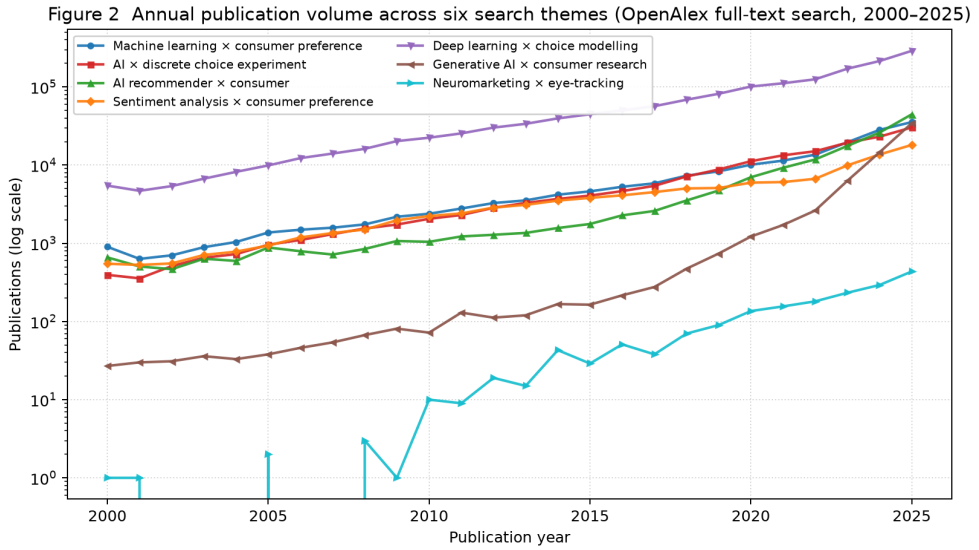


**Figure 2.** Annual publication volume across six themes (OpenAlex full-text search, 2000–2025, log scale)

The six curves fall into three diffusion patterns. Sustained exponential: machine learning and recommender themes keep growing near-exponentially, indicating infrastructure-like status. Event-driven: generative AI stayed low until 2022 then grew an order of magnitude in two years, matching the public release of large language models. Steady growth: sentiment analysis, neuromarketing and AI × DCE grow smoothly, the first benefiting from review data accumulation, the latter two constrained by survey cost and methodological entry barriers. These patterns imply different forecasting risks: exponential series are most sensitive to functional form, event-driven series are nearly unforecastable from history, and steady series are the most reliably fitted.

### *4.2. Growth-curve fits and extrapolation*

Table 2 reports growth estimates. For machine learning × consumer preference, Logistic gives $K = 1065480.0$, $b = 0.23296$, inflection 2039.69, $R^2 = 0.97523$; Gompertz gives $K = 1065480.0$, $b = 0.05322$, inflection 2048.51, $R^2 = 0.95452$. For AI × discrete choice experiment, Logistic gives $K = 910020.0$, $b = 0.19991$, inflection 2042.07, $R^2 = 0.99548$; Gompertz gives $K = 910020.0$, $b = 0.04566$, inflection 2052.33, $R^2 = 0.98565$. All $R^2$ exceed 0.95, yet $K$ is far above 2025 values and inflections mostly fall beyond the sample, indicating growth has not saturated. Figure 3 shows fits and 2026–2030 extrapolation.

**Table 2.** Growth-curve parameter estimates and extrapolation (asymptotic SE in parentheses)

| Series | Model | $K$ | $b$ | Inflection | $R^2$ | 2026 | 2030 |
|---|---|---|---|---|---|---|---|
| ML × consumer preference | Logistic | 1065480.0 (10108790.7) | 0.23296 (0.03694) | 2039.69 | 0.97523 | 42129.4 | 100835.8 |
| ML × consumer preference | Gompertz | 1065480.0 (3629117.41) | 0.05322 (0.04262) | 2048.51 | 0.95452 | 38742.1 | 73143.1 |
| AI × discrete choice experiment | Logistic | 910020.0 (3615592.59) | 0.19991 (0.01329) | 2042.07 | 0.99548 | 35210.8 | 74789.6 |
| AI × discrete choice experiment | Gompertz | 910020.0 (1693345.81) | 0.04566 (0.01999) | 2052.33 | 0.98565 | 32631.0 | 56874.2 |

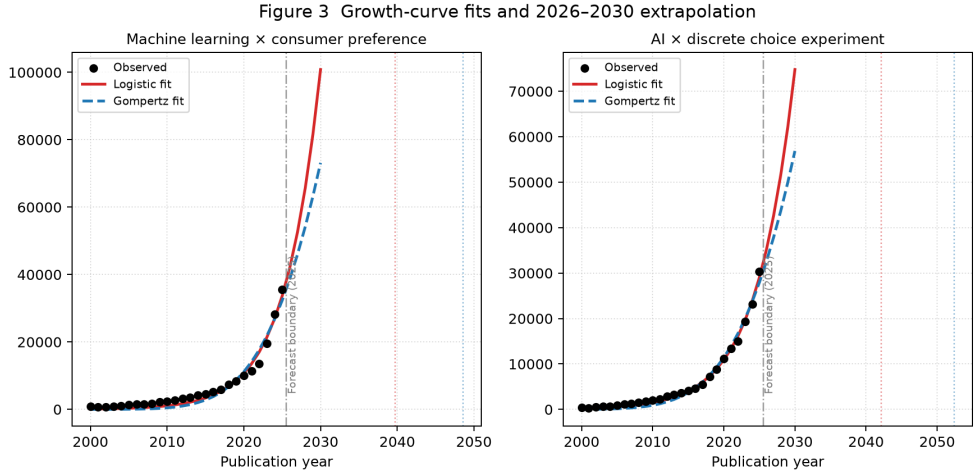


**Figure 3.** Growth-curve fits and 2026–2030 extrapolation (left: ML × consumer preference; right: AI × DCE)

Inflections fall around 2040 (Logistic) and 2048 (Gompertz) for machine learning, and 2042/2052 for AI × DCE; estimated $K$ (roughly 0.91–1.07 million) far exceeds current levels, placing both series in a take-off/acceleration phase. The models disagree on extrapolation: Gompertz assumes slower pre-inflection acceleration and a longer post-inflection tail, yielding lower 2030 values than Logistic (e.g. ~73k versus ~101k for the machine-learning theme). Such divergence is itself evidence of long-run uncertainty; we treat extrapolations as scenarios rather than point forecasts.

### *4.3. Country, subject and document-type distributions*

Of 24,954 title-and-abstract records (2010–2025), the United States leads with 4,552 papers, followed by India (3,417), China (1,529) and the United Kingdom (1,460), with Germany, Australia and Canada next (Figure 4, left). By subfield, marketing leads with 4,490, followed by artificial intelligence (2,791), information systems (1,863) and electrical/electronic engineering (1,757), forming a "marketing-problem, computing-method" interdisciplinary structure (Figure 4, right). By document type, journal articles (8,798), book chapters (8,563) and conference papers (4,978) dominate, with preprints (928) and dissertations (705) reflecting rapid iteration.

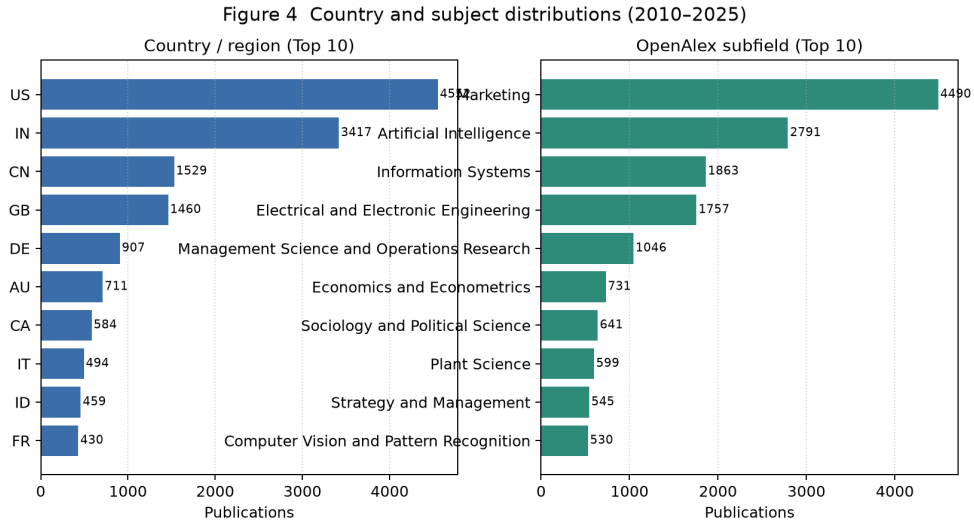


**Figure 4.** Country (left) and OpenAlex subfield (right) distributions, 2010–2025

**Table 3.** Country and subject distributions (Top 10, 2010–2025)

| Country | Count | Subfield | Count |
|---|---|---|---|
| US | 4552 | Marketing | 4490 |
| IN | 3417 | Artificial Intelligence | 2791 |
| CN | 1529 | Information Systems | 1863 |
| GB | 1460 | Electrical and Electronic Engineering | 1757 |
| DE | 907 | Management Science and Operations Research | 1046 |
| AU | 711 | Economics and Econometrics | 731 |
| CA | 584 | Sociology and Political Science | 641 |
| IT | 494 | Plant Science | 599 |
| ID | 459 | Strategy and Management | 545 |
| FR | 430 | Computer Vision and Pattern Recognition | 530 |

A long tail is also evident, and while China and India are rising, high-impact work remains concentrated in a few English-speaking institutions. This cumulative advantage has substantive external-validity implications: if most AI preference models are trained and validated on WEIRD (Western, educated, industrialised, rich, democratic) samples, generalisation across cultures and institutions requires care[47]. Generative AI and online survey platforms lower the marginal cost of cross-cultural data collection, offering a route to correct this bias—although the cultural validity of synthetic respondents remains unproven[18].

### *4.4. Method-family frequencies*

Aggregated totals (2000–2025; Figure 5, Table 4) put AI-augmented choice modelling first at about 1.74 million (two queries), followed by ML segmentation (178,649), recommendation (144,898), NLP sentiment (107,242), generative AI (63,724) and visual/behavioural data (1,817). The choice-

modelling total is inflated by the broad "deep learning × choice modelling" full-text query, which includes papers merely mentioning deep learning, so conceptual dilution is a caveat. Generative AI has a recent-decade share of 98.1%, concentrating almost all output in the last ten years; recommendation and sentiment have shares of 89.4% and 73.9%. The structure reflects both scaled application of mature methods and explosive growth of new ones.

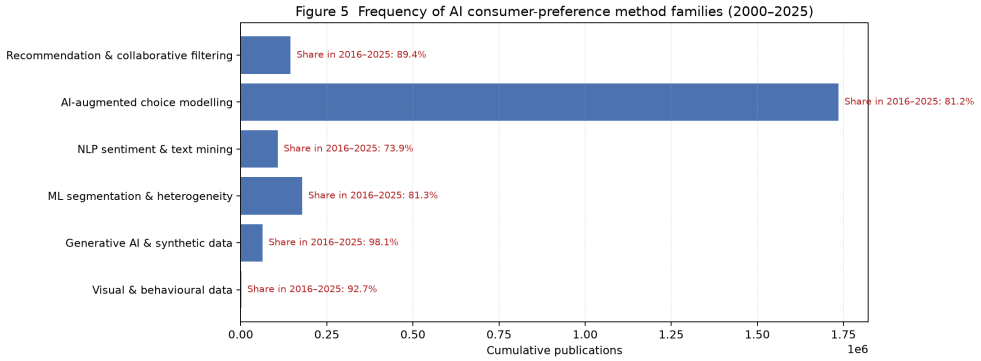


**Figure 5.** Cumulative publications by method family and recent-decade share

**Table 4.** Method-family frequency statistics (2000–2025)

| Method family | Total 2000–2025 | Total 2016–2025 | Recent share (%) |
|---|---|---|---|
| Recommendation & collaborative filtering | 144898 | 129478 | 89.4 |
| AI-augmented choice modelling | 1735631 | 1408714 | 81.2 |
| NLP sentiment & text mining | 107242 | 79225 | 73.9 |
| ML segmentation & heterogeneity | 178649 | 145315 | 81.3 |
| Generative AI & synthetic data | 63724 | 62516 | 98.1 |
| Visual & behavioural data | 1817 | 1684 | 92.7 |

The frequency structure also reveals two opportunities. Visual and behavioural data is the smallest family (1,817) but has the highest recent share (92.7%), a small fast-growing field: eye tracking and neural signals provide process information unavailable from choice data, but high equipment cost and small samples limit generalisability; combining physiological measures with wearables and mobile logs is a promising direction. Recommendation and sentiment mining are mature (89.4% and 73.9% recent shares), implying rich benchmarks but limited algorithmic novelty; incremental innovation lies more in integrating them with structural preference models. Frequency statistics thus indicate not merely scale but the lifecycle stage of each family.

### *4.5. Negative binomial regression*

Table 5 reports the paper-level negative binomial model. Sample $n = 40$, mean citations 132.9, median 30.5, open access 62.5%, journal articles 62.5%. Pearson dispersion is 288.206, far above 1, with overdispersion $\alpha = 2.09583$; the Poisson–NB likelihood-ratio test is $\chi^2 = 6044.814$ ($p = 0.0$), favouring NB. Publication year has IRR 0.7312 ($p = 0.00045$), i.e. each additional year changes conditional mean citations by about -26.9%. Open access IRR is 1.8816, journal article IRR 1.9702; with marketing as reference, AI/CS IRR is 2.0268 and business/economics/other 1.034. Pseudo $R^2 = 0.96935$,

but this is inflated with tiny samples and strong overdispersion and is not a reliable fit measure. Only publication year is significant at 5% and negative—likely reflecting that "best match" retrieval privileges older highly cited papers, not a causal "newer means fewer citations" effect. Open access ($p = 0.24403$), article type ($p = 0.19287$) and subject dummies are not significant. With $n = 40$, these coefficients describe sample associations only, not causal effects.

**Table 5.** Negative binomial regression of citations (robust SE; IRR = incidence rate ratio)

| Variable | Coef. | Robust SE | $z$ | $p$ | IRR | IRR 95% CI |
|---|---|---|---|---|---|---|
| Constant | 5.02091 | 0.79716 | 6.299 | 0.0 | 151.5493 | [31.7688,722.949] |
| Year (centred) | -0.31302 | 0.08918 | -3.51 | 0.00045 | 0.7312 | [0.614,0.8709] |
| Journal article | 0.67815 | 0.52079 | 1.302 | 0.19287 | 1.9702 | [0.7099,5.4679] |
| Open access | 0.63213 | 0.54261 | 1.165 | 0.24403 | 1.8816 | [0.6496,5.4501] |
| AI/Computer Science | 0.70647 | 0.70262 | 1.005 | 0.31466 | 2.0268 | [0.5114,8.0334] |
| Business/Econ/Other | 0.03343 | 0.58668 | 0.057 | 0.95455 | 1.034 | [0.3274,3.2652] |
| $\alpha$ | 2.09583 Pearson dispersion=288.206 pseudo $R^2$=0.96935 $n = 40$ | | | | | |

Three methodological cautions follow. Citations are age-dependent, so any citation model containing publication year must control for it. The distribution is extremely right-skewed (median 30.5, mean 132.9, SD 285.4), explaining the large $\alpha$ and making mean-based inference fragile. Open access and article type are often treated as impact drivers, but their effects are confounded with journal prestige, disciplinary norms and research quality; with limited samples they are better treated as controls. The value of this section lies in demonstrating that citation data require count models and dispersion diagnostics, not in substantive claims about the determinants of impact.

### *4.6. Comparative-evidence matrix*

Table 6 summarises verifiable comparisons. The overall picture is conditional advantage rather than wholesale substitution: deep learning beats logit/mixed logit only with samples $\geq O(10^4)$, high-dimensional inputs or complex mechanisms[21]; TasteNet-MNL exceeds benchmark MNL and mixed logit on Swissmetro and recovers nonlinear taste functions[31]; but in street-view image choice, DCMs are nearly equal in out-of-sample prediction and more interpretable[22]. In heterogeneity modelling, G-MNL-type models often outperform MXL and LCM, yet no model dominates[9,26]. For generative AI, LLM synthetic respondents explain just 1.3% of individual and 4.2% of country-level human variance, with 77.8% recovery of demographic gradients but only 48.1% of cultural dimensions[18]; factorial-survey digital twins also compress response variance and inflate normativity. These numbers provide testable benchmarks for the discussion.

**Table 6.** Verifiable evidence matrix from method comparisons

| Family | Comparison | Source | Quantitative finding | Implication |
|---|---|---|---|---|
| Choice modelling | DNN vs MNL/MXL predictive comparison | Wang, Wang, Bailey & Zhao (2021), Transportation Research Part B 148:60-81 | With sample size $\geq O(10^4)$, high-dimensional inputs or a complex true data-generating process, DNNs outperform BNL, BXL, MNL and MXL in zero-one loss, log loss and approximation loss; under the opposite conditions they perform worse. | Provides boundary conditions for AI-augmented choice models |
| Choice modelling | TasteNet-MNL vs MXL comparison | Han, Pereira, Ben-Akiva & Zegras (2022), Transportation Research Part C | On the Swissmetro dataset TasteNet-MNL achieves higher predictive accuracy than benchmark MNL and mixed logit and recovers nonlinear taste functions; misspecified logit gives biased parameters. | Neural learning of systematic taste heterogeneity improves prediction |
| Choice modelling | DCM vs DNN/CNN in image-based choice | Zhu & Si (2024), Journal of Choice Modelling (street-view cycling-choice DCE) | Discrete choice models show nearly equal out-of-sample predictability to DNNs and significantly outperform CNNs; DCMs are more explainable; semantic-interpretation data outperform raw imagery. | AI does not necessarily beat classical DCM; data representation matters |
| Heterogeneity modelling | Comparison of six heterogeneity models | Keane & Wasi (2013), Journal of Applied Econometrics 28(6):1018-1045 | N-MIXL and latent class are not dominant; G-MNL, MM-MNL and T-MIXL typically outperform them, yet no model dominates universally. | RPL and LCM are complementary; choose by research purpose |
| Heterogeneity modelling | MXL/LC/LC-MXL comparison of PEV preferences | Jia & Chen (2023), Transportation Research Part A (Virginia PEV DCE) | All three agree that monetary incentives are most effective, charging infrastructure second and battery-range gains weak; no model is unanimously superior across fit, prediction and behavioural interpretation. | Multi-model triangulation beats any single model |
| Heterogeneity modelling | GMXL vs LCM comparison in wine choice | Goncalves et al. (2019), Economics Letters 178:28-32 (wine choice) | Generalised mixed logit accounting for both taste and scale heterogeneity fits best (higher McFadden pseudo-$R^2$, lower AIC/BIC, higher correct-prediction rate); LCM is more useful for market segmentation. | Best fit and segmentation usefulness are distinct goals |
| Generative AI | Global audit of LLM synthetic respondents | Chen, Wang & Bhatia (2026), Nature Human Behaviour (74-country virtual GPS survey) | LLMs explain only 1.3% of individual-level and 4.2% of country-level human variance; demographic gradients are recovered in 77.8% of cases but cultural dimensions in only 48.1%; MSE $\approx$ 1.827. | Synthetic respondents cannot yet replace representative human surveys |
| Generative AI | LLM factorial-survey digital twins | EPC2026 conference paper (China n=5,186; US n=5,906 factorial family-ideals survey) | LLM response variance is compressed (1.1–2.1 vs human 2.4–2.6), with normative inflation and flattened subgroup differences; the gendered response pattern cannot be recovered. | Uncalibrated synthetic responses overstate consensus and understate contestation |
| Text mining | Aspect-based sentiment + House of Quality framework | Ray & Singh (2025), Journal of Retailing and Consumer Services (smartwatch reviews) | InstructABSA-based aspect sentiment mining combined with HoQ and PROMETHEE-II constructs a relative performance index across 12 smartwatch brands and generates recommendations. | Review sentiment can be converted into preference weights for MCDM |
| Recommender systems | Meta-analysis of online recommendation agents | Journal of Retailing (2023), meta-analysis of online recommendation agents | Recommendation agents have significant but context-moderated effects on consumer decisions; personalisation transparency and trust are key boundary conditions. | Commercial impact of recommenders is context-dependent |

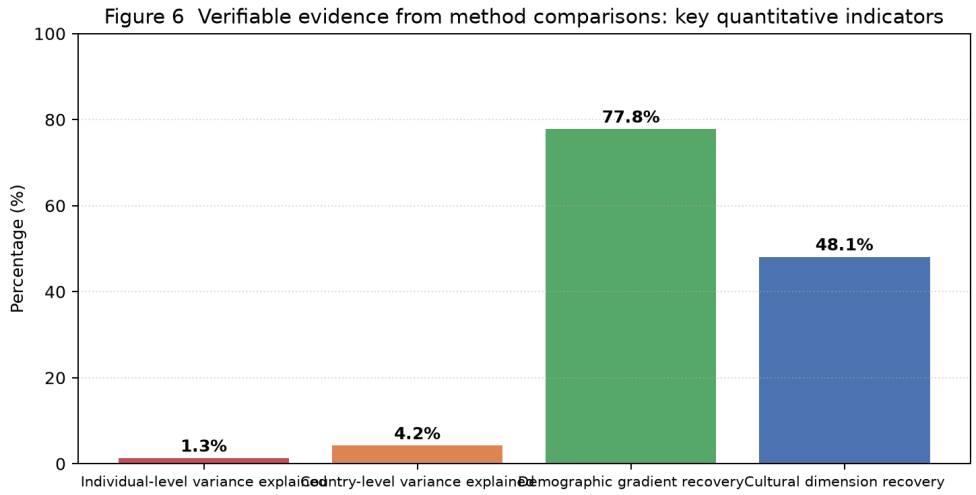


**Figure 6.** Key quantitative indicators from the global audit of LLM synthetic respondents (Chen et al., 2026)

Three implications follow. First, evidence on deep learning versus choice models does not support "AI wins": predictive advantage is conditional, and gains often trade against interpretability, consistent with statistical learning theory on estimation–approximation error trade-offs[20,21]. Second, the "no free lunch" result in heterogeneity modelling means model choice should follow research purpose rather than fit statistics: LCM's intuitive segmentation, RPL's continuous welfare distribution and G-MNL's joint treatment of heterogeneity each suit different decisions[9,26]. Third, the synthetic-respondent evidence is the most cautionary: even when group-level regularities are reproduced, individual and cultural validity remain weak, and systematic deviations may mask genuine social disagreement[18,39]. Together these support our "complementarity, not substitution" stance.

## 5. Discussion

### *5.1. Strengths and boundaries of each method family*

Recommender systems capture "what users will choose" from large interaction logs, excelling at prediction and scale, but cold start, sparsity and filter bubbles are inherent limits, and their logic rarely explains "why"[13,47]. AI-augmented choice modelling preserves random-utility interpretability while increasing functional flexibility, but gains depend on sample size and dimensionality, and it does not outperform classical models in small samples[21,30]. NLP sentiment mining expands data sources by turning reviews into structured preference evidence, but aspect-identification errors, irony and fake reviews introduce measurement bias[16,51]. ML segmentation detects latent groups in high-dimensional interactions but needs XAI tools to become actionable marketing insight[37,49]. Generative AI offers data augmentation and scenario generation, but current synthetic respondents cannot replace representative samples[18]. Visual and behavioural data capture implicit preferences but are costly, privacy-sensitive and of limited ecological validity.

Researchers must also weigh three kinds of cost. Data cost: DCE requires design, pre-testing and large surveys but yields controlled quality; recommendation and review mining reuse platform logs at low marginal cost but suffer selection bias (only active users leave traces); eye tracking and EEG are expensive with small samples, suiting mechanism studies rather than population inference. Ex-

pertise cost: mixed logit and LCM require econometric training, deep models require engineering, and their fusion requires cross-disciplinary teams. Ethical and compliance cost: the finer the behavioural data, the higher the privacy and consent requirements; synthetic data alleviates privacy pressure but introduces validity risk. Researchers should trade off under budget, capability and ethical constraints rather than chase technical fashion. This helps explain why, despite its leading frequency, AI-augmented choice modelling has produced successful deployments mainly where large data and algorithmic teams coexist.

### *5.2. Complementarity with DCE, RPL and LCM*

Our core argument is complementarity rather than substitution. First, specification discovery: neural-learned nonlinear utility functions can supply data-driven priors for mixed-logit coefficient distributions, reducing subjective specification search[31]. Second, heterogeneity characterisation: LCM classes and ML clusters can cross-validate, the former offering probabilistic membership and statistical tests, the latter complex boundaries in high-dimensional feature space[27,28]. Third, data fusion: review sentiment, eye tracking and clickstreams can enter choice models as auxiliary variables, alleviating state dependence and unobserved heterogeneity bias[32]. Concrete hybrids exist: Sifringer et al. embed a neural network as a flexible utility function inside a discrete-choice framework, retaining standard logit probability structure[30]; TasteNet-MNL delegates the individual-characteristics-to-taste mapping to a network while constraining behavioural indicators such as value of time and elasticities to plausible ranges[31]; aspect-level sentiment scores (e.g. battery life, after-sales service) can enter choice models as attribute-level covariates to explain unobserved heterogeneity[32]. In such designs AI handles representation learning and pattern discovery, while econometric models handle probability structure, identification and behavioural interpretation—preserving credibility for counterfactual policy simulation.

A practical decision map follows from arranging families along prediction–explanation. With large interaction logs and short-term prediction goals, recommendation and collaborative filtering are preferred; with attribute-level willingness to pay and moderate samples, DCE + RPL/LCM remains the gold standard; with unstructured text, NLP sentiment converts reviews into preference evidence; with high-dimensional hidden segments, ML clustering and probabilistic LCM classification complement each other; with data scarcity or privacy constraints, generative AI supports scenario generation and pre-testing but should not replace final sampling; with interest in choice-formation processes, visual and behavioural data provide unique process evidence. The underlying logic again reflects Breiman's multiplicity of statistical cultures and Shmueli's separation of explanation and prediction[19,20].

### *5.3. Privacy, interpretability and synthetic-data validity*

AI preference analysis depends on fine-grained behavioural data, creating three governance challenges. Privacy: clickstreams, locations and physiological signals are highly identifiable; de-identification does not remove re-identification risk via data linkage, requiring differential privacy, federated learning and management norms[47]. Interpretability: deep models' accuracy often comes with opaque decision processes; users and regulators cannot see why a recommendation or price was set; XAI and model cards help but the "fidelity of explanations" remains contested[38,48,50]. Synthetic data: LLM digital twins reproduce some group-level demographic gradients but system-

atically compress variance and flatten subgroup differences; uncalibrated use in policy evaluation may overstate consensus and understate controversy[17,18]. Synthetic data are therefore better suited to pre-testing, hypothesis generation and privacy-sensitive exploration than to final inference. We recommend a three-layer "data–model–application" review framework. At the data layer, examine legality and representativeness: informed consent for behavioural data, assessment of re-identification risk beyond de-identification, and reporting of key demographic distributions to judge generalisability. At the model layer, examine faithfulness and robustness: explanation attributions should match true model mechanisms, verified through adversarial testing, sensitivity analysis and subgroup stability; for generative AI, report benchmarking of synthetic versus real data in distribution, correlation structure and group differences, not merely surface similarity. At the application layer, examine consequences and accountability: when models inform pricing, recommendation or policy simulation, assess differential impacts on vulnerable groups and establish human review and appeal channels. This framework is not intended to block innovation but to build accumulated trust—without trust, even high predictive accuracy cannot translate into sustainable commercial and social value[48,49].

## 6. Conclusion and Research Agenda

Taking "AI as a method and tool for analysing consumer preferences" as our object, and combining bibliometrics, growth and negative binomial models, and verifiable comparative evidence, we draw three conclusions. First, AI consumer-preference research expanded by orders of magnitude over 2000–2025, with generative AI exploding in the last three years, yet extrapolation suggests most themes have not saturated and the field remains in rapid flux. Second, the six method families have distinct boundaries: AI's predictive advantage presupposes large samples, high-dimensional inputs and complex mechanisms; in small-sample and interpretation-intensive settings, traditional DCE, mixed logit and latent class remain indispensable, and multi-model triangulation beats any single model. Third, privacy, interpretability and synthetic-data validity form a "governance triangle" that must be embedded in methodological innovation.

Our marginal contributions are threefold. Object-wise, rather than taking a single industry or algorithm as the unit, we place six families in one framework, avoiding industry-specific confounding of method evaluation. Evidentially, we splice macro-level trends, paper-level correlations and peer-reviewed comparison figures across three layers, so claims about evolution are supported by scale data and individual verifiable numbers. Normatively, we reject the linear "AI replaces traditional models" narrative and advance complementarity and governance, with actionable agendas. These contributions are not final: limited by OpenAlex scope and micro-sample size, our growth extrapolations and citation model should be read as methodological demonstrations and testable hypotheses rather than established empirical laws.

We propose four research agendas. First, hybrid architectures: systematically compare combinations such as neural utility specification + mixed logit estimation, LCM + cluster validation, and sentiment-augmented choice models across sample sizes and attribute dimensionalities, establishing contingency rules for method choice. Second, validity boundaries of synthetic data: test LLM synthetic respondents' distributional, structural and policy-extrapolation validity within cross-cultural DCE frameworks, developing standard protocols with variance calibration and human-anchored priors. Third, interpretable and responsible preference inference: introduce SHAP, counterfactual

explanations and fairness constraints into preference models and study their causal impact on consumer trust and adoption[37,49]. Fourth, data governance and measurement norms: establish ethical norms for collecting, de-identifying and using visual/behavioural data, and develop methods that preserve preference heterogeneity while protecting privacy[47]. Their realisation requires infrastructure and evaluation upgrades: cross-cultural benchmark datasets and reproducible protocols, multi-objective evaluation beyond prediction accuracy (behavioural plausibility, external validity, interpretability, computational cost), and collaboration among econometricians, computer scientists and privacy researchers. Through such agendas, AI can evolve from a tool for predicting preferences into a partner for understanding them, enhancing decision efficiency while safeguarding behavioural explanation and value rationality.

**Author Contributions:** Conceptualization, Z.L. and X.T.; methodology, Z.L.; software, Z.L.; validation, Z.L. and X.T.; formal analysis, Z.L.; investigation, Z.L.; resources, Z.L.; data curation, Z.L.; writing—original draft preparation, Z.L.; writing—review and editing, X.T.; visualization, Z.L.; supervision, X.T.; project administration, Z.L. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The data supporting the findings are available from the corresponding author upon reasonable request.

**Acknowledgments:** The authors thank the editors and reviewers for their constructive comments.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## References

1. Simon, H. A. A behavioral model of rational choice. *Quarterly Journal of Economics* **1955**, *69(1)*, 99-118.
2. Kahneman, D.; Tversky, A. Prospect theory: an analysis of decision under risk. *Econometrica* **1979**, *47(2)*, 263-291.
3. Mcfadden, D. Conditional logit analysis of qualitative choice behavior. ZAREMBKA P. Frontiers in Econometrics. New York: Academic Press, 1974; pp. 105-142.
4. Mcfadden, D. Econometric models of probabilistic choice. MANSKI C F, MCFADDEN D. Structural Analysis of Discrete Data with Econometric Applications. Cambridge, MA: MIT Press, 1981; pp. 198-272.
5. Mcfadden, D.; Train, K. E. Mixed MNL models for discrete response. *Journal of Applied Econometrics* **2000**, *15(5)*, 447-470.
6. Train, K. E. Discrete choice methods with simulation. 2nd ed. Cambridge: Cambridge University Press, 2009.
7. Greene, W. H.; Hensher, D. A. A latent class model for discrete choice analysis: contrasts with mixed logit. *Transportation Research Part B: Methodological* **2003**, *37(8)*, 681-698.
8. Hensher, D. A.; Greene, W. H. The mixed logit model: the state of practice. *Transportation* **2003**, *30(2)*, 133-176.
9. Keane, M.; Wasi, N. Comparing alternative models of heterogeneity in consumer choice behavior. *Journal of Applied Econometrics* **2013**, *28(6)*, 1018-1045.
10. Mayer-Schonberger, V.; Cukier, K. Big data: a revolution that will transform how we live, work, and think. Boston: Houghton Mifflin Harcourt, 2013.

11. Lecun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521(7553)*, 436-444.
12. Goodfellow, I.; Bengio, Y.; Courville, A. Deep learning. Cambridge, MA: MIT Press, 2016.
13. Ricci, F.; Rokach, L.; Shapira, B. Recommender systems handbook. 2nd ed. New York: Springer, 2015.
14. Koren, Y.; Bell, R.; Volinsky, C. Matrix factorization techniques for recommender systems. *Computer* **2009**, *42(8)*, 30-37.
15. Hu, Y.; Liu, X. Mining and summarizing customer reviews. Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York: ACM, 2004; pp. 168-177.
16. Pang, B.; Lee, L. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* **2008**, *2(1-2)*, 1-135.
17. Argyle, L. P.; Busey, T. J.; Gubler, J. R.; et al. Out of one, many: using language models to simulate human samples. *Political Analysis* **2023**, *31(3)*, 337-351.
18. Chen, F.; Wang, Q.; Bhatia, S. Large language models as synthetic respondents for economic preferences: a global audit. *Nature Human Behaviour* **2026**.
19. Breiman, L. Statistical modeling: the two cultures. *Statistical Science* **2001**, *16(3)*, 199-231.
20. Shmueli, G. To explain or to predict?. *Statistical Science* **2010**, *25(3)*, 289-310.
21. Wang, S.; Wang, Q.; Bailey, N.; Zhao, J. Deep neural networks for choice analysis: a statistical learning theory perspective. *Transportation Research Part B: Methodological* **2021**, *148*, 60-81.
22. Zhu, W.; Si, W. Predicting choices of street-view images: a comparison between discrete choice models and machine learning models. *Journal of Choice Modelling* **2024**.
23. Jia, W.; Chen, T. D. Investigating heterogeneous preferences for plug-in electric vehicles: policy implications from different choice models. *Transportation Research Part A: Policy and Practice* **2023**.
24. Louviere, J. J.; Hensher, D. A.; Swait, J. D. Stated choice methods: analysis and application. Cambridge: Cambridge University Press, 2000.
25. Ben-Akiva, M.; Lerman, S. R. Discrete choice analysis: theory and application to travel demand. Cambridge, MA: MIT Press, 1985.
26. Goncalves, T.; Costa, Pinto L. M.; Lourenco-Gomes, L. Exploring distinct sources of heterogeneity in discrete choice experiment: an application to wine choice across European consumers. *Economics Letters* **2019**, *178*, 28-32.
27. Vermunt, J. K.; Magidson, J. Latent class cluster analysis. HAGENAARS J A, VAN DER HEIJDEN P G M, MUTHEN B O. Applied Latent Class Analysis. Cambridge: Cambridge University Press, 2002; pp. 89-106.
28. Wedel, M.; Kamakura, W. A. Market segmentation: conceptual and methodological foundations. 2nd ed. Boston: Kluwer Academic, 2000.
29. Zhang, S.; Yao, L.; Sun, A.; Tay, Y. Deep learning based recommender system: a survey and new perspectives. *ACM Computing Surveys* **2019**, *52(1)*, Article 5.
30. Sifringer, L.; Lurkin, V.; Alahi, A. Integrating machine learning into discrete choice models: the neural network as a flexible utility function. *Transportation Research Part C: Emerging Technologies* **2020**, *111*, 559-580.
31. Han, Y.; Pereira, F. C.; Ben-Akiva, M.; Zegras, C. A neural-embedded discrete choice model: learning taste representation with strengthened interpretability. *Transportation Research Part C: Emerging Technologies* **2022**, *136*, 103538.
32. Ray, R. K.; Singh, A. From online reviews to smartwatch recommendation: an integrated aspect-based sentiment analysis framework. *Journal of Retailing and Consumer Services* **2025**, *82*, 104059.
33. Devlin, J.; Chang, M. W.; Lee, K.; Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT. 2019: 4171-4186.
34. Vaswani, A.; Shazeer, N.; Parmar, N.; et al. Attention is all you need. Advances in Neural Information Processing Systems. 2017, 30: 5998-6008.
35. Breiman, L. Random forests. *Machine Learning* **2001**, *45(1)*, 5-32.

36. Chen, T.; Guestrin, C. XGBoost: a scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York: ACM, 2016; pp. 785-794.
37. Lundberg, S. M.; Lee, S. I. A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems. 2017, 30: 4765-4774.
38. Ribeiro, M. T.; Singh, S.; Guestrin, C. Why should I trust you?: explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York: ACM, 2016; pp. 1135-1144.
39. Bisbee, J.; Larregue, D.; Nyhan, B.; et al. Synthetic replacements for human survey respondents? A study of large language models in political polling. *Political Analysis* **2024**, *32(4)*, 401-416.
40. Turing, A. M. Computing machinery and intelligence. *Mind* **1950**, *59(236)*, 433-460.
41. Samuel, A. L. Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development* **1959**, *3(3)*, 210-229.
42. Haykin, S. Neural networks and learning machines. 3rd ed. Upper Saddle River: Pearson, 2009.
43. Schmidhuber, J. Deep learning in neural networks: an overview. *Neural Networks* **2015**, *61*, 85-117.
44. Krizhevsky, A.; Sutskever, I.; Hinton, G. E. ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems. 2012, 25: 1097-1105.
45. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016: 770-778.
46. Sutton, R. S.; Barto, A. G. Reinforcement learning: an introduction. 2nd ed. Cambridge, MA: MIT Press, 2018.
47. Boyd, D.; Crawford, K. Critical questions for big data. *Information, Communication and Society* **2012**, *15(5)*, 662-679.
48. Pasquale, F. The black box society: the secret algorithms that control money and information. Cambridge, MA: Harvard University Press, 2015.
49. Guidotti, R.; Monreale, A.; Rugieri, F.; et al. A survey of methods for explaining black box models. *ACM Computing Surveys* **2018**, *51(5)*, Article 93.
50. Doshi-Velez, F.; Kim, B. Towards a rigorous science of interpretable machine learning. arXiv:1702.08608, 2017.
51. Liu, B. Sentiment analysis and opinion mining. San Rafael: Morgan and Claypool, 2012.


**Disclaimer/Publisher's Note:** All statements, views and data in published papers belong to the named authors and contributors, not to Lightup Innovation or the editors. Neither Lightup Innovation nor the editors accept liability for loss or damage arising from ideas, methods, instructions or products mentioned in the content.