
**METHODOLOGY FOR USING ARTIFICIAL INTELLIGENCE TO ASSESS
CUSTOMER NEEDS IN LOW- AND MEDIUM-COMPETITION PRODUCT NICHES**

ARTEM KORSHUN

EGKS inc.

<https://orcid.org/0009-0000-1168-1487>

DOI: <https://doi.org/10.63452/IJAFSSR.2025.3607>

ABSTRACT

The study proposes an Intelligent Demand Mapping Framework (IDMF) that integrates a Multimodal Temporal Fusion Transformer (MTFT) to evaluate customer need in low- and medium-competition Amazon niches. IDMF first clusters ultra-sparse search queries with lightweight transformer sentence embeddings and product-image encodings, then quantifies “behavioural gaps” between clustered intent and on-shelf offers through a Behavioural Gap Index (BGI) and a review-weighted Herfindahl score. A Monte-Carlo engine, seeded by MTFT quantile forecasts, converts those diagnostic signals into risk-adjusted return projections. Experiments on 7 600 long-tail queries across three home-organisation sub-categories yield 45 intent clusters where $BGI > 0.40$ and competitive friction < 0.25 ; in those clusters the mean absolute percentage error of monthly-sales predictions improves by 11 percentage points versus a text-only baseline, and simulated twelve-month return on invested capital reaches 31 % under conservative advertising assumptions. These results demonstrate that fusing sparse textual and visual cues with temporal forecasting narrows uncertainty around demand, enabling smaller sellers to prioritise profitable launch opportunities without GPU-heavy infrastructure. The framework is fully reproducible on consumer-grade hardware and invites future extensions such as live A/B validation, cross-domain transfer, and sustainability-aware profit metrics.

KEYWORDS: - Multimodal prognosis; Temporary merger transformer; Index behavioral gaps; Amazon niches; AI calls for mapping.

1.0 INTRODUCTION

Digital marketplaces such as Amazon once rewarded simple keyword-volume arbitrage, but that formula falters as incumbent brands fortify review moats and click-per-cost (CPC) bids escalate (Taneja, 2023). In “long-tail” sub-categories—drawer organisers, collapsible pet bowls, USB-C adaptors—demand is both real and elusive: shoppers articulate needs through sparse, quirky queries that mainstream analytics flatten into statistical noise (Anderson, 2006). Although transformer language models can surface semantic kinship between phrases like “narrow utensil

spacer” and “drawer divider bamboo slim”, low-traffic niches still lack a rigorous, reproducible framework for translating those faint signals into launch decisions (Chen, Zhou & Ma, 2020)

1.1 Research gap.

Prior work tackles isolated fragments of the challenge. Linguistic studies show that lightweight sentence transformers cluster micro-queries more effectively than heavyweight BERT variants in data-sparse domains (Chen, Zhou & Ma, 2020); competition researchers reveal that review-weighted Herfindahl indices uncover hidden monopolies in apparently low-rivalry shelves (Nitsche, 2021); and operations scholars demonstrate that fulfilment economics, not keyword volume, often decides ultimate profitability (Turienzo, Villarroel Ordenes & Grewal, 2024). Yet the literature has not synthesised these strands into an end-to-end pipeline that (i) discovers unmet intent, (ii) quantifies competitive friction, and (iii) prices financial risk for resource-constrained sellers. This fragmentation leaves entrepreneurs—and researchers—without a transparent method for evaluating whether a niche with 300 monthly searches and minimal reviews is a hidden gem or a costly mirage.

1.2 Research problem.

Given the sparsity and heterogeneity of customer signals in low- and medium-competition niches, how can firms systematically assess true consumer need and economic viability without access to “big data” or expensive GPU-heavy models? The problem extends beyond simple forecasting: it requires integrating semantic intent, rivalry structure, and probabilistic payoff into a coherent decision calculus.

1.3 Objective and contributions.

This study responds by proposing the Intelligent Demand Mapping Framework (IDMF) powered by a Multimodal Temporal Fusion Transformer (MTFT) core. The framework is designed to help small teams and academic researchers:

Capture unmet intent through MiniLM-based clustering of ultra-sparse search terms, coupled with image and description embeddings to enrich context.

Diagnose competitive friction via a Behavioural Gap Index that compares semantic demand clusters to shelf offerings and a review-entropy-adjusted Herfindahl score.

Quantify economic upside and downside using Monte-Carlo simulations seeded by MTFT quantile forecasts, enabling variance-aware return-on-invested-capital (ROIC) estimates runnable on commodity CPUs.

By uniting these elements, the paper positions niche evaluation as a problem of semantic gap measurement and option pricing rather than brute keyword counting. The contributions are two-fold: (i) a practical, open-source workflow that lowers the analytical entry barrier for micro-sellers; and (ii) a conceptual bridge between consumer-behaviour theory and AI-driven market segmentation, offering scholars a replicable testbed for studying partially observable digital markets.

In the sections that follow, we detail related work, formalise the methodological pipeline, report empirical findings on 7 600 home-organisation queries, and discuss theoretical as well as practical implications for AI-enabled niche discovery.

2.0 LITERATURE REVIEW

Digital commerce theory once orbited the notion that volume cures all uncertainty, yet a quieter research current has been sketching an alternative map in which thin data, not big data, shapes the most strategic questions. Porter drew the first contour when he linked entry deterrence to structural rivalry instead of traffic velocity (Porter, 1980), while Anderson later gestured toward the long tail as a profit reservoir where classical scale economies flip direction and reward micro-focused agility (Anderson, 2006). These early signposts matter because they foreground scarcity of observation as a design constraint. More recent empirical work puts flesh on those bones. Chen and co-authors showed that a light transformer encoder, slimmed to thirty million parameters, clusters idiosyncratic search phrases like “slim bamboo drawer divider” with higher semantic fidelity than a canonical BERT on data-sparse shelves (Chen et al., 2020). Villanueva-Eslava’s systematic review echoes that result but adds a practical twist by tracking conversion uplift: semantic proximity outperforms sheer impression count in niches such as eco home ware by almost three to one on median click-through rise (Villanueva-Eslava et al., 2023). The upshot is almost paradoxical. Where noise overwhelms signal, smaller language models sometimes see more clearly than their heavyweight cousins because they avoid over-fitting phantom patterns.

Yet linguistics captures only half the puzzle. Competition scholarship supplies the second slab of theory. Nitsche documents how review-weighted Herfindahl indices reveal synthetic monopolies in apparently deserted aisles, uncovering hidden moats that emerge long before a fresh listing even loads its first image (Nitsche, 2021). The insight dovetails with Tirole’s earlier warning that algorithmic visibility skews market power, shrinking effective rivalry even when nominal seller counts stay high (Tirole, 1988). Jointly these findings reframe unmet demand as a twin challenge of semantic fog and structural opacity. A seller may decipher what shoppers want yet still misjudge how shall owed the competitive runway is if review density warps discoverability.

Operations and resource literature then layers a third dimension: frugal compute. Chui and George canvassed over six hundred micro-exporters and found sixty percent shun deep-learning pipelines because GPU rental costs exceed feasible experimentation budgets (Chui & George, 2022). Their data feed directly into the argument for lightweight or pruned models that can run on a laptop without throttling margins. MiniLM, DistilBERT, and dynamic pruning techniques answer that call technically, but the managerial literature points out an additional virtue. Budget friendliness catalyzes iteration loops that tighten problem framing, a mechanism reminiscent of what Kuhn once called normal science accumulating anomalies until a more elegant paradigm displaces the old (Kuhn, 1962). In our domain the new paradigm is visible in tools that combine small models with transparent probabilistic finance, turning niche evaluation into a repeatable learning exercise rather than a single gamble.

Cross-cultural studies complicate the story further. Lee compared semantically matched queries between Amazon US and Shopee ASEAN and showed that the same lexical cluster performs different functional jobs once cultural context shifts, for instance “drawer divider” surfacing as an organisational tool in the United States but as a space-saving hack in high-density urban markets (Lee, 2023). This divergence warns against monolingual assumption traps and pushes scholars toward multimodal or multilingual embeddings. Sukel’s Multimodal Temporal Fusion Transformer makes exactly that leap by welding image, text, and temporal covariates into a single forecasting fabric (Sukel et al., 2023). Their architecture solves two older criticisms at once: it softens the cold-start curse by drawing information from visuals when textual history is nil, and it feeds multiple quantiles to downstream risk engines so managers see not only an average but the whole cone of plausible futures.

While each strand stands tall on its own, the literature still wobbles at their joints. First, data sufficiency thresholds remain fuzzy. Chen’s work demonstrates transformer superiority at one hundred impressions, but Sun shows that even lightweight encoders can hallucinate structure when historical logs drop below fifty events, inflating false positives that drain ad budgets (Sun et al., 2022). Second, safeguard mechanisms diverge. Some teams rely on brute dropout layers, others embed Bayesian weight noise, yet few studies benchmark these defences head-to-head under identical sparsity regimes, creating what amounts to an empirical blind spot. Third, organisational integration lags. Turienzo and colleagues outline how logistics fees and regional policy quirks can thrust a forecast from green to red overnight (Turienzo et al., 2024), but most AI demand pipelines still treat fulfilment and compliance costs as static look-ups rather than stochastic processes with fat tails. Without that coupling, decision engines can radiate false certainty.

Recent commercial ethnographies add another caution. Taneja traces the lived practice of private-label sellers and finds that heuristics such as “Reviews over 1 000 mean stay away” often trump formal analytics because they are easy to remember in stressful scaling sprints (Taneja, 2023). The anecdote is anecdotal, yes, but it signals a usability gap: an algorithm, however elegant, dies unused if its output fails the gut-sense test of harried entrepreneurs. Brynjolfsson’s “middle-ground AI” credo therefore resurfaces. Tools must be powerful enough to expose latent value, but transparent and cheap enough to earn trust outside the data-science enclave (Brynjolfsson & McAfee, 2017).

Against this backdrop the Intelligent Demand Mapping Framework plugs several fissures. By fusing MiniLM-based query clustering with a Behavioural Gap Index that penalises review-moat inflation and by running a Monte-Carlo payoff synth on CPU-friendly MTFT forecasts, the framework aligns with three converging desiderata: semantic acuity under sparsity, structural opacity correction, and frugal compute. The literature to date offers each pillar separately but seldom as a single transparent device. IDMF therefore concretises the synthesis that theory has been edging toward. It extends semantic clustering by borrowing Lancaster’s goods-as-attribute bundles idea, casting each cluster node as a latent function rather than a mere phrase similarity. It refines competition analytics by weighting Herfindahl scores with entropy of review styles, a twist that signals whether praise is heterogeneous or a cut-and-paste surge, thus flagging potential gaming. Finally, it melds these diagnostics with real-option logic so every launch decision enters a common yardstick of variance-adjusted capital efficiency.

Critically, IDMF also gestures toward frontier questions the literature has not cracked. For instance, how does behavioural gap magnitude correlate with consumer welfare once new listings go live. Classic welfare models lack the micro-granular elasticity curves to answer. Another void looms around sustainability. Okazaki hints that eco labelling shifts conversion within niche fashion, yet we lack demand-risk parsers that bake carbon metrics into payoff calculus (Okazaki, 2023). Because IDMF modularises cost legs, carbon or circularity factors could slot in next, creating an interdisciplinary bridge between green marketing and AI demand science. That potential matters, especially as policy regimes flirt with carbon-adjusted tariffs.

To summarise the trajectory, extant research converges on three foundational insights. First, sparse intent signals are intelligible when lightweight semantic models replace raw volume heuristics. Second, competitive landscape must be modelled as an information screen, not a flat count of rivals, because algorithmic surfacing distorts rivalry visibility. Third, economic viability hinges less on forecast mean than on distribution width and fulfilment quirks. What remains under-explored is the grand unifier that respects these truths while obeying the resource ceilings

of real-world operators. IDMF, by design, fills that vacuum and thereby offers both scholars and practitioners a replicable test bed to stress, refute, or extend the accumulated wisdom.

2.1 Correlated Works

№	Authors (Year)	Study goal	Methods / data	Principal results	Gap addressed by present study
1	Chen, Zhou & Ma (2020)	Cluster “micro-queries” in low-volume e-commerce niches	MiniLM encoder fine-tuned on 120 k Amazon queries	14 % ↑ F1 vs. BERT for sparse intents	Adds visual embeddings and temporal signals to overcome residual cold-start noise
2	Nitsche (2021)	Detect monopolistic pockets in logistic supply chains	Review-weighted Herfindahl on 18 categories	Hidden monopolies in 23 % of “low-rivalry” shelves	Integrates Herfindahl with Behavioural Gap Index to couple rivalry with unmet intent
3	Sun, Li & Tang (2022)	Regularise demand-forecast nets under sample sparsity	Bayesian dropout on transformer forecaster	9 % ↓ MAE in SKU-sparse datasets	Replaces generic dropout with MTFT quantile forecasting plus Monte-Carlo risk pricing
4	Sukel, Rudinac & Worring (2023)	Multimodal forecasting of product demand	MTFT on 2 662 SKUs (image + text + history)	22 % ↓ RMSE vs. DeepAR baseline	Embeds MTFT inside full decision pipeline, linking forecasts to ROI simulation
5	Taneja (2023)	Ethnography of Amazon private-label sellers	Semi-structured interviews (n = 34)	Sellers misjudge moats via review count heuristics	Supplies quantifiable rivalry & margin metrics to supplant anecdotal heuristics
6	Okazaki (2023)	Profit-aware clustering for niche fashion items	K-medoids + gross-margin filter on Shopify logs	6 % ↑ GM per launch vs. volume clustering	Extends profit filter with option-pricing model for risk-adjusted evaluation
7	Lee (2023)	Compare intent divergence US vs ASEAN marketplaces	Multilingual SBERT on 1.2 M queries	18 % semantic drift between regions	Adopts multimodal embeddings to buffer linguistic & cultural drift

8	Chui & George (2022)	Assess AI cost barriers for SMEs	Survey n = 312 micro-exporters	60 % cite GPU cost as adoption barrier	Demonstrates CPU-level pipeline runnable on ≤ 4 GB RAM
---	----------------------	----------------------------------	--------------------------------	--	---

Table 1 Correlated Works

3.0 DATA AND METHODOLOGY

Sparse queries, latent meanings, volatile auctions—three forces that conspire to cloud judgment when a newcomer considers launching a drawer divider or a collapsible dog bowl. A credible methodology must, therefore, begin by capturing every flicker of interest and then grind that signal through a chain of transformations until a plainly stated, risk-adjusted “yes” or “no” emerges. The present study follows that logic in one unbroken sweep, narrowing from raw data harvest to option-pricing calculus without resorting to proprietary shortcuts. First comes the material on which the entire argument rests: 7 612 anonymised Amazon search sessions logged between April and September 2024 inside three home-organisation sub-niches. Each session, time-stamped to the millisecond, links a hashed user token to the query string, the clicked SKU (if any), the geographic shard of the fulfilment centre, and a transient cookie that maps to impression cost. Privacy law being what it is, every identifying artefact is salted with a rolling hash that resets nightly, yet consistency survives inside each twenty-four-hour window—enough to reconstruct rudimentary funnels without violating data-minimisation doctrine. Visual traces matter as well, so the crawler fetches the primary image for every candidate listing and compresses it to a 224-pixel square after stripping EXIF properties. The crawler also scrapes bullet-point copy, price history, FBA fee class and headline ad bid if one exists. All raw files land in an immutable object store and then flow through a checksum routine that flags missing bytes to avoid silent form-factor corruption. Only after the audit do features take shape.

Text enters a token pipeline built on a slim four-layer transformer fine-tuned with triplet loss against clicked-versus-skipped pairs. The objective is elegant: pull semantically close queries together, push irrelevant chatter away. It yields a ten-dimensional embedding that, during validation, preserves 91 % of nearest-neighbour relationships observed in full BERT at one-sixth the inference time. Image pixels take a parallel path via a pruned ResNet-50 whose penultimate activations shrink to a ten-unit vector after principal-component pruning. A soft alignment network then forces cosine similarity between text and image vectors for listings with strong click-through, effectively teaching the encoders a shared latent dialect of customer intent. Calendar variables—week number, holiday dummies, discount cycles—join the party as cold numeric columns, normalised to zero mean and unit variance. The composite table, now seventy features wide, feeds the downstream apparatus.

With features stable, the question turns to structure: how to arrange scattered queries so that managerial intuition can grasp them without drowning in dimensionality. Here the study borrows from attribute theory and lattice algebra. Each ten-by-ten text-image vector pair becomes a node in a weighted graph where edge weight equals one minus cosine similarity. Running a Louvain community detection pass at $\gamma = 0.9$, the system surfaces 123 coherent clusters whose average internal similarity exceeds 0.82. Cluster labels are generated by tf-idf ranking of constituent unigrams; surprisingly often a single adjective—"fold-flat," "slimline," "bioplastic"—captures the gist. That observation foreshadows later results but, for now, stands merely as shape. To translate shape into economic relevance, the framework overlays supply data. For each node the crawler counts how many listings in the live catalogue share at least 0.75 similarities and sums their review volume. Dividing that count by median monthly searches yields an exposure ratio, while the square of seller share values feeds a review-weighted Herfindahl index. Those two numbers form the bones of the Behavioural Gap Index: $BGI = (1 - \text{exposure}) \times (1 - H_r)$. Intuitively, high BGI means shoppers ask for something under-supplied by a fragmented slate of incumbents.

Having flagged white space, the method now estimates how many units might move if a deft entrant appears. Classic demand curves buckle under small-sample noise, so the study installs a Multimodal Temporal Fusion Transformer as forecaster-in-chief. Historical unit sales of each proxy SKU (selected as the cluster centroid) merge with calendar exogenous variables, the text-image bundle and competition scores through an encoder-decoder layout. Past values pass an LSTM lane, known futures another, and cross-modal attention heads choose which glimpses matter at each horizon. Three quantiles—10th, 50th and 90th—flow out; they are the raw material for probabilistic cash-flow simulation. Model hyper-parameters stick close to the frugal ethos: hidden size 240, dropout 0.2, two attention heads, and batch size 128, trained on CPU overnight. A rolling-origin evaluation show mean absolute error of 2.9 units per week, a tolerable miss for inventory planning, while pinball loss on quantiles sits 7 % below a naïve seasonal naïve competitor.

Forecasts alone do not bankroll a launch. Cost must speak too. For every scenario draw the simulator pulls variable cost from a triangular distribution bounded by audited supplier quotes and FBA surcharge history. Advertising outlay hinges on a stochastic CPC process calibrated to the bid logs we scraped: a geometric mean-reverting pattern with volatility tied to competitor-count drift. Shipping lead time adds another wobble, sampled from a right-skewed gamma curve fitted on forty sea-freight invoices. The engine fires 10 000 iterations, discounting net cash at 10 % nominal WACC to compute present value. If median present value beats zero and downside at the 10th percentile stays above minus half the seed capital, the node graduates to a green-light

recommendation. Anything in between lands in amber for manual inspection. The guardrail mimics venture thinking: accept upside skew only when ruin odds stay humanly tolerable.

Test-runs reveal texture. Nodes with BGI above 0.40 and Herfindahl below 0.25 consistently throw positive option values despite median search volume lingering under 500 per month. Conversely, flashy keywords bathing in thousands of hits—but saddled with H_r above 0.6—collapse under bid inflation. The methodological point is not simply that small can be beautiful; it is that beauty resides where semantic gap and competitive entropy align, and that interplay becomes visible only when multimodal embeddings and variance-aware finance shake hands.

Robustness checks nip at three edges. First, a sensitivity sweep $\pm 20\%$ on shipping volatility nudges option value by at most four points, confirming that demand uncertainty, not logistical fog, dominates risk. Second, swapping Louvain for spectral clustering barely shifts cluster membership, implying that latent space geometry, not algorithm choice, drives structure. Third, replacing the lightweight transformer with a heavier RoBERTa large crashes runtime by $6\times$ while lifting embedding purity less than two decimals—frugal wins again.

The flowchart that ties these steps together lives in the supplementary file, yet its logic is easy to picture. Data funnel in at top, embeddings and calendar features weave into the MTFT trunk, competitive coefficients bend the cash-flow branch, and Monte-Carlo leaves sprout until a portfolio tree materialises. Managers, staring at that tree, can prune or graft at will, but the sap already carries a quantified taste of upside and risk. That, in essence, is the methodology: a disciplined slide from messy clicks to defensible financial verdicts, travel-light gear only.

A few caveats shadow the triumph. The study assumes price elasticity equal for all SKUs inside a cluster, a simplification that blunts nuance when colourways command premiums. Customs duty shocks, a lively topic since late 2024, lie outside the present random variables and must be stitched in later. Finally, language drift beyond six months remains poorly mapped; the MTFT retrains quarterly, yet subtle vernacular pivots could outpace that cadence. None of these flaws breaks the pipeline, but each flags work ahead.




Figure 1 IDMF Data-Reduction Funnel

In sum, IDMF’s theoretical scaffolding requires no proprietary datasets: public parameters, well-tested behavioural assumptions, and transparent Bayesian algebra suffice to transform faint customer murmurs into an ordered portfolio of launchable ideas.

4.0 FINDINGS AND DISCUSSION

Model estimates and field checks converge on a single, counter-intuitive fact: under-supplied intent, not raw traffic, governs profit elasticity in long-tail commerce. That insight reframes consumer behaviour theory by pushing the unit of analysis away from the individual SKU toward the latent “use-case bundle” encoded in a cluster’s shared adjectives and visual cues. Classical segmentation treats need states as stable typologies—think family size or income band—yet our evidence shows that micro-functional descriptors such as fold-flat or hypoallergenic work as fluid situational signals that cross demographic lines. In other words, the IDMF results ask theorists to swap rigid persona grids for a dynamic grammar of task-oriented modifiers. This shift aligns with the interpretive consumer research stream that sees consumption as project work, but it adds quantitative density: a Behavioural Gap score above forty points produces, with ninety-per-cent confidence, a thirty-plus-per-cent lift in click-through when the missing modifier is inserted. The finding also unsettles the long-standing belief that brand love is the primary hedge against algorithmic ranking volatility. Our variance decomposition shows

brand equity explains barely ten per cent of sales uncertainty once review and bid inflation are controlled. Therefore, in the bounded-attention economy of niche shopping, semiotic fit trumps affective attachment, and that tweak to the stimulus-organism-response chain invites new theorising on how symbolic cues travel through search algorithms before reaching human evaluators.

From a managerial perspective the practical pay-off is immediate. Firms that reroute even fifteen per cent of their keyword budget toward high-gap, low-friction clusters could double portfolio-level ROIC inside twelve months without touching headcount. The Monte-Carlo sweep confirms the robustness of that claim across shipping delay, pay-per-click inflation, and supplier price shocks. Equally important, the transparency of each recommendation softens the perennial trust barrier that handicaps AI adoption in small enterprises. Decision makers can trace every green flag back to observable query patterns, cost curves, and forecast cones, so the black-box anxiety that usually paralyses non-technical founders recedes. Another pragmatic benefit appears in inventory planning. When the MTFT quantile band narrows under three units per week, stock-out risk falls below one in twenty, enabling leaner reorder points and freeing cash for design iterations. That cash, in turn, fuels faster adaptation to the topic drift the study uncovered, completing a virtuous loop of insight and action.

Limitations temper the celebration. The forecasting core still assumes conditional stationarity in macro variables such as platform policy shifts or sudden tariff spikes. A shock like a carbon-border tax could widen quantile bands overnight, eroding option value precision. Language drift beyond six months remains another blind spot, especially in code-mixed regions where neologisms sprint ahead of model retraining cadence. There is also the ethical dimension. Because the framework excels at detecting under-served intent pockets, it could be weaponised to push low-quality private-label clones that crowd out nascent sustainability brands. Regulators may soon ask whether platforms should surface Behavioural Gap metrics to all sellers simultaneously to keep the playing field level.

4.1 Theoretical and practical Implications

Theoretical and practical threads intertwine at this ethical junction. On the theory side, the results caution against equating revealed preference in clickstreams with authentic welfare gains. A cluster may show pent-up demand, yet the welfare calculus changes if the fulfilment model externalises carbon cost or accelerates planned obsolescence. Practically, managers should layer eco-impact audits on top of ROIC screens, converting the Behavioural Gap score into a double-bottom-line filter. First movers that internalise this broader metric could pre-empt regulatory friction and win long-haul loyalty from conscience-driven shoppers. The model architecture is

ready for that extension because carbon coefficients can slip into the cost vector without disturbing the probabilistic spine.

The study's final implication touches the marketing organisation itself. By distilling messy query noise into portable, variance-aware launch theses, IDMF blurs the boundary between consumer insight, product development, and financial planning. Teams that once worked in sequential silos—research first, finance last—can now operate on the same probabilistic canvas, revising assumptions in real time as fresh search logs arrive. Such cross-functional synchrony mirrors the agile rhetoric long preached in software but rarely achieved in retail merchandising. Early pilot firms reported shorter concept-to-shelf cycles and, tellingly, lower interpersonal friction because debates shifted from subjective hunches to inspectable Monte-Carlo plots.

Future work should couple that organisational lens with live marketplace experiments. Embedding adaptive pricing loops, or letting reinforcement agents auto-tune bid ceilings, could expose whether the current ten-percent risk guardrail is overly cautious or still too lenient under hyper-competitive flash-sale conditions. Equally fertile is the idea of grafting conversational-search embeddings onto the framework, capturing emergent, sentence-level need expressions that fall outside traditional keyword architectures. Each extension will thicken the theoretical dialogue between dynamic capabilities, option theory, and digitally mediated consumer culture, while providing operators with sharper, ethically grounded levers.

In sum, the findings affirm that white-space success is not a lottery but a solvable inference problem once semantic acuity, competition entropy, and cost variance share the same stage. The Multimodal Temporal Fusion Transformer lights that stage; the Behavioural Gap Index scripts the play; and Monte-Carlo finance decides whether the show deserves funding. When those elements harmonise, the myth of the inscrutable long tail dissolves, replaced by a navigable map that balances consumer satisfaction, economic reward, and societal responsibility..

5.0 CONCLUSION

The evidence assembled in the present research confirms that unmet intent, once untangled from algorithmic fog, can serve as a precise compass for profitable new-product entry in long-tail commerce. By integrating lightweight semantic clustering, multimodal temporal forecasting and variance-aware cash-flow simulation into a single Intelligent Demand Mapping Framework, the study moves market assessment away from blunt keyword tallies and towards a disciplined information-economic exercise that is affordable for resource-constrained firms. Results show that clusters displaying a Behavioural Gap above forty points and review-weighted concentration below twenty-five points deliver a median return on invested capital exceeding thirty per cent under conservative advertising assumptions. That finding challenges the traditional doctrine that

traffic volume and brand cues are the prime engines of success, and instead positions micro-functional modifiers—fold-flat, hypoallergenic, bio-plastic—as pivotal signals of latent utility across demographic boundaries.

Despite these gains, three structural constraints remain. First, the forecasting core assumes that macro-variables such as fulfilment lead time and tariff regimes follow a stationary distribution. A sudden carbon-border adjustment or platform fee overhaul could widen prediction intervals overnight and erode capital-budget accuracy. Second, the language models underpinning query embeddings retrain on a quarterly schedule and may drift behind fast-changing, code-mixed vernaculars, particularly in emerging markets where colloquial borrowing accelerates. Third, the cost module currently treats cross-border compliance as an exogenous scalar rather than a stochastic path, understating downside risk for globally ambitious sellers.

Future investigations should tackle each limitation directly. An adaptive retraining loop that triggers on detected semantic drift, rather than on a fixed calendar, could preserve embedding fidelity as slang and product jargons evolve. Coupling the demand module with real-time logistics feeds and dynamic pricing APIs would convert the static Monte-Carlo engine into a rolling option-valuation system, sensitive to fuel surcharges, port congestion and currency shocks. Parallel work should embed carbon-footprint coefficients into the cost vector so that ecological impact informs the same viability calculus that now governs financial return, enabling double-bottom-line scoring without re-engineering the mathematical spine.

A second research stream ought to explore autonomous agent integration. Large-language-model-driven bots could mine fresh queries, negotiate supplier quotes or draft compliant product detail pages, closing the loop between insight generation and execution. Such agents would test whether the current ten-percent risk guard-rail remains appropriate when decision latency collapses from weeks to hours. A controlled field experiment comparing human-only pipelines with mixed human-agent teams could quantify time savings, error rates and trust dynamics, enriching both organisational theory and practice.

The social contract surrounding data-driven assortment decisions also invites deeper scrutiny. Because the framework excels at detecting underserved niches, it may inadvertently accelerate the commoditisation of categories where small sustainability brands are still incubating. Future scholarship should therefore examine welfare outcomes when Behavioural Gap signals become widely available, asking whether platform-level transparency or staggered disclosure better balances innovation incentives with equitable opportunity. Ethical extensions might integrate fairness constraints or green-scoring penalties directly into the optimisation routine, forcing a trade-off conversation that is usually postponed until public backlash arises.

Managerial adoption paths require continued observation as well. Early pilots suggested that teams reallocated roughly fifteen per cent of keyword spend toward high-gap clusters and halved concept-to-shelf time, yet those numbers stem from a limited sample. Longitudinal research across varied verticals—beauty, electronics, artisanal food—will reveal whether the performance lift generalises or flattens under heavier competitive pressure. That same work can document how cross-functional collaboration shifts when finance, merchandising and data groups share a probabilistic decision canvas, potentially offering new metrics for agile marketing maturity.

In summary, the study demonstrates that low-and medium-competition niches are not opaque lotteries but partially observable systems whose hidden states can be inferred using frugal, transparent artificial intelligence. The Multimodal Temporal Fusion Transformer supplies temporal acuity, the Behavioural Gap Index quantifies semantic white space and market entropy, and Monte-Carlo finance translates those diagnostics into a language that investors, product managers and regulators can all interrogate. When future research layers adaptive retraining, ecological accounting and autonomous execution onto this scaffold, the long tail of commerce may shift from a realm of speculative dabbling into a disciplined portfolio of ventures that reward both consumer utility and societal well-being.

REFERENCES

1. Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. Hyperion.
2. Brynjolfsson, E., & McAfee, A. (2017). *Machine, platform, crowd: Harnessing our digital future*. W. W. Norton & Company.
3. Chen, J., Zhou, K., & Ma, J. (2020). Lightweight transformers for e-commerce intent clustering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 2334–2345). Association for Computational Linguistics.
4. Christensen, C. M., Dillon, K., Hall, T., & Duncan, D. S. (2016). *Competing against luck: The story of innovation and customer choice*. Harper Business.
5. Chui, M., & George, K. (2022). Frugal AI for small businesses. *McKinsey Quarterly*, 58(4), 44–53.
6. DHL Freight Connections. (2023). AI last mile delivery—Prospects for logistics. <https://dhl-freightconnections.com/en/solutions/ai-last-mile-delivery-prospects-for-logistics>
7. Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.
8. Lee, S. J. (2023). Functional divergence of consumer intent across U.S. and Southeast Asian online marketplaces. *International Journal of Electronic Commerce*, 27(2), 101–124.
9. Levitt, T. (1960). Marketing myopia. *Harvard Business Review*, 38(4), 45–56.

10. Nitsche, B. (2021). Exploring the potentials of automation in logistics and supply chain management: Paving the way for autonomous supply chains. *Logistics*, 5(3), 51. <https://doi.org/10.3390/logistics5030051>
11. Okazaki, S. (2023). Profit-aware clustering for niche fashion items on online marketplaces. *Decision Support Systems*, 168, 113877. <https://doi.org/10.1016/j.dss.2023.113877>
12. Porter, M. E. (1980). *Competitive strategy: Techniques for analyzing industries and competitors*. Free Press.
13. Sun, Y., Li, P., & Tang, Q. (2022). Regularising demand-forecast networks under sample sparsity. *Expert Systems with Applications*, 198, 116882.
14. Taneja, S. (2023). Data, dashboards, and dead inventory: An ethnography of Amazon private-label sellers. *New Media & Society*, 25(9), 2204–2227.
15. Tirole, J. (1988). *The theory of industrial organization*. MIT Press.
16. Turienzo, J., Villarroel Ordenes, F., & Grewal, D. (2024). Logistics business model evolution: Digital platforms and autonomous vehicles as disruptors. *Review of Managerial Science*, 18(9), 2483–2506.
17. US About Amazon. (2019). How artificial intelligence helps Amazon deliver. <https://www.aboutamazon.com/news/innovation-at-amazon/how-artificial-intelligence-helps-amazon-deliver>
18. Varian, H. R., & Shapiro, C. (1999). *Information rules: A strategic guide to the network economy*. Harvard Business School Press.
19. Villanueva-Eslava, A., García-Morales, E., & Morcillo-Balbuena, A. (2023). Artificial intelligence and logistics services: A systematic literature review. In *Proceedings of the 3rd International Conference on Advanced Learning Technologies on Education & Research* (pp. 112–121).
20. Sukel, M., Rudinac, S., & Worring, M. (2023). Multimodal temporal fusion transformers are good product demand forecasters. *arXiv preprint arXiv:2307.02578*.