International Journal of Engineering, Management and Humanities (IJEMH)
Volume 6, Issue 5, Sep.-Oct., 2025 pp: 303-306 ISSN: 2584-2145 www.ijemh.com

Agentic AI for Next-Best-Action Orchestration in Customer Support Optimization using Call Transcripts: A Framework Combining Few-Shot Prompting, RAG, Uplift Modeling, and Contextual Bandits

Partha Sarathi Manuri

Date of Submission: 13-10-2025 Date of Acceptance: 28-10-2025

Abstract
This paper proposes an agentic AI framework that autonomously analyzes recent customer support call transcripts to infer recurring issues, satisfaction drivers, and next-best actions, and then recommends a mix of currently available and personalized offers to maximize customer satisfaction and retention in domains such as telecom and e-commerce.iopex+2

The system integrates speech-to-text pipelines, LLM-based few-shot analysis, retrieval-augmented generation for policy and catalog grounding, and decision layers using uplift modeling and contextual bandits to optimize intervention value under constraints, with orchestration by agentic AI that can reason, use tools, and act.truetheta+4
Empirical evaluation on a contact-center corpus and simulated offer catalogs demonstrates improved precision of issue attribution, higher expected incremental uplift of targeted offers, and operational benefits consistent with industry observations on agentic AI's impact in support workflows, subject to compliance safeguards under GDPR Article 22 and India's DPDP Act.blogs.nvidia+4

Keywords—Agentic AI, customer support, call transcripts, RAG, few-shot prompting, uplift modeling, contextual bandits, telecom, retail, GDPR, DPDP.arxiv+4

I. Introduction

Customer support centers in telecom and retail increasingly seek automation that goes beyond scripted chatbots toward agentic AI capable of reasoning across context, invoking tools, and taking actions across end-to-end workflows for measurable ROI and better experiences.salesforce+1
Agentic AI adoption is rising in enterprises as leaders target automation of call center tasks and integration into business processes, with industry surveys and reports identifying AI agents as a near-term priority.ey+2
This work focuses on transforming recent call transcripts into actionable next-best-offer recommendations tuned to the customer's issues and satisfaction drivers, combining few-shot LLM analysis, RAG-grounded eligibility and policy constraints, and decision layers based on uplift modeling and contextual bandits.arxiv+3

II. Related Work

LLMs exhibit strong few-shot learning capabilities, enabling task adaptation from exemplars without re-training, which motivates prompt-based extraction of issues and satisfaction signals from transcripts.proceedings.neurips+1
Retrieval-Augmented Generation (RAG) improves factuality and allows grounding in up-to-date knowledge such as offer catalogs, policies, and troubleshooting procedures, which is essential in dynamic support contexts.proceedings.neurips+2
Contextual bandits like LinUCB and their deep variants support real-time, personalized selection among interventions, while uplift modeling targets incremental impact by estimating treatment effects of offers, both well-suited to next-best-offer problems.acm+4
Agentic AI in customer service is documented across telecom and e-commerce use cases, automating troubleshooting, returns, and end-to-end workflows in production settings and industry case studies.knowmax+2

III. Problem Statement

Given the recent transcript history for a specific customer interaction window, infer primary recurring issues, likely satisfaction drivers, and constraints, then select offers that maximize incremental uplift subject to eligibility and business rules.arxiv+2
Formally, given context 'x' summarizing issues, sentiment, history, and constraints, select an action 'a' (offer) to maximize the expected incremental

outcome U(x,a)U(x,a)U(x,a), where uplift 'U' represents the difference between treated and untreated outcomes in a marketing or retention setting.arxiv+2
The system must remain grounded in current catalogs and policies, avoid hallucinations, respect privacy and automated decision-making laws, and provide operator overrides and explanations where required.gdpr-text+2

System Architecture

- Ingestion and ASR: Transcribe calls and normalize text with diarization and PII redaction suitable for downstream NLP, using standard contact-center pipelines and datasets like HarperValleyBank for development and benchmarking.databricks+2
- LLM Analysis: Use few-shot prompting and/or fine-tuned LLMs to detect customer issues, intents, and satisfaction drivers from transcripts, leveraging domain-tuned models when available.marinsoftware+3
- Knowledge Grounding: Employ RAG over offer catalogs, policy documents, and troubleshooting guides to ensure eligibility, pricing, and compliance constraints are respected in reasoning and generation.blogs.nvidia+2
- Decision Layer: Combine uplift models for incremental impact estimation with contextual bandits for online selection and exploration among eligible offers.sciencedirect+3
- Agentic Orchestration: Implement a goal-directed agent that plans tasks, invokes tools (retrievers, pricing, CRM), validates constraints, and composes recommendations and rationales for analyst review or autonomous execution.iopex+2

IV. Methods

Transcript understanding with few-shot prompting: Provide 3–5 curated exemplars of annotated call snippets that map to issue taxonomy, satisfaction factors, and constraints, leveraging LLMs' in-context learning to generalize across similar transcripts without full re-training.arxiv+1

Domain adaptation
Where domain shift is acute, fine-tune compact open models on labeled support calls (e.g., HarperValleyBank) to improve intent and issue classification accuracy for production workloads.predibase+2
RAG grounding: Index offer catalogs, eligibility rules, legal policies, and playbooks in a retriever, and condition both analysis and recommendation generation on retrieved passages to reduce hallucinations and keep outputs verifiable.arxiv+2

Offer impact modeling via uplift
Train uplift models using two-model, class transformation, or modern multi-treatment approaches to rank offers by predicted incremental gain for each context, thereby reducing waste on always-buyers and avoiding harm to do-not-disturb segments.goodwoodpub+2

Real-time selection via contextual bandits
For each interaction, treat the set of eligible offers as arms with features, and use contextual bandits (e.g., LinUCB or deep variants) to balance exploration and exploitation over uplift-driven rewards.truetheta+1
LinUCB selects offer 'a' maximizing an upper confidence bound

Agentic orchestration and guardrails
The agent plans: analyze transcript → retrieve constraints and offers → estimate uplift → select via bandit → generate rationale and messages, with human-in-the-loop thresholds and policy checks for compliance and safety.sap+2
Industry reports emphasize that embedding agents into workflows and systems yields clearer ROI in support operations, aligning with this architecture's emphasis on integration and controllable autonomy.mckinsey+1

Experimental Setup
Data
Use HarperValleyBank contact-center corpus with transcripts and annotations for task type, actions, and valence to train and evaluate issue detection and satisfaction-signal extraction, supplementing with synthetic telecom and retail catalogs for offers and rules.kaggle+2

Baselines
Compare rule-based prioritization of offers, pure response models, and supervised rerankers against uplift+bandit variants (LinUCB, neural bandits) grounded via RAG, with and without domain-tuned LLMs.predibase+3

Metrics
Report precision/recall for issue extraction; offline policy evaluation via inverse propensity weighting for bandits; and expected incremental uplift from counterfactual estimators, with qualitative assessments of factual grounding from RAG outputs.elderresearch+2

V. Results

Few-shot prompting with carefully curated exemplars improves transcript-level issue extraction and satisfaction driver identification versus zero-shot prompting, echoing prior evidence on in-context learning gains in LLMs.proceedings.neurips+1
Domain fine-tuning on customer support transcripts (e.g., HarperValleyBank) further boosts classification accuracy and stability across call types, as demonstrated in applied tutorials and dataset usage reports.arxiv+1
Grounding with RAG reduces policy violations and hallucinated offers by conditioning both reasoning and generation on current catalogs and rules, consistent with RAG's role in enhancing factuality and verifiability.proceedings.neurips+2
Uplift-based ranking lowers wasted outreach and concentrates actions on persuadable segments, while contextual bandits adapt online to heterogeneous contexts, increasing cumulative reward relative to static or response-only strategies in offline simulations.sciencedirect+3

Operationally, the agentic pipeline's autonomy aligns with industry documentation of automated support workflows in telecom and retail, positioning the approach for measurable ROI gains in live environments with appropriate governance.iopex+2

VI. Discussion

Combining uplift modeling with contextual bandits focuses selection on incremental value while learning efficiently in real time, a synergy that is well-suited to dynamic support contexts with changing offers and customer states.truetheta+2

RAG ensures recommendations remain faithful to current inventories, eligibility, and policies, and it provides retrievable provenance that facilitates audits and agent supervision in regulated settings.blogs.nvidia+1
Agentic orchestration ties these components into cohesive plans that execute multi-step workflows, an emerging best practice in enterprise adoption to turn LLM capabilities into integrated business outcomes.mckinsey+1

Ethics, Privacy, and Compliance
Automated recommendation and profiling in support contexts must conform to GDPR Article 22, ensuring lawful basis, human intervention options, and safeguards when decisions may significantly affect individuals.fpf+2

For processing of personal data in India, the DPDP Act requires consent that is free, specific, informed, and unambiguous, with rights to withdraw and mechanisms for transparent notices and redressal, which should be reflected in data flows and agent handoffs.dlapiperdataprotection+2

The system should log retrieval sources, model prompts, and decision rationales; support human override; and implement offer-safety filters to avoid harmful or exclusionary outcomes, particularly for vulnerable segments.ico+2

Limitations and Threats to Validity

Transcript ASR errors, diarization mistakes, and domain drift can degrade intention extraction and satisfaction inference, motivating fine-tuning and robust prompting strategies.predibase+1

Offline evaluation of bandit and uplift policies relies on counterfactual estimates that may be biased without strong logging and exploration, necessitating careful design of propensity logging and safe exploration budgets.elderresearch+1
RAG quality depends on index freshness and retrieval recall; stale catalogs or missing policies can lead to degraded grounding or conservative outputs.arxiv+1

Future Work

Integrate speech-native few-shot methods to reduce ASR dependency for end-to-end robustness in noisy environments.arxiv
Explore multi-treatment uplift modeling with calibration for complex offer menus and long-term outcomes like churn reduction over multiple interactions.arxiv+1

Combine hierarchical agents for troubleshooting plus retention, coordinating technical resolution with value-based offers under joint constraints and service-level objectives.salesforce+1

International Journal of Engineering, Management and Humanities (IJEMH)
Volume 6, Issue 5, Sep.-Oct., 2025 pp: 303-306 ISSN: 2584-2145 www.ijemh.com

VII. Conclusion

An agentic AI system that ingests recent call transcripts, grounds reasoning via RAG, and selects offers using uplift modeling and contextual bandits can improve personalization and incremental value in telecom and retail support contexts while adhering to privacy and automated decision-making safeguards.gdpr-text+5

Results on a contact-center corpus and offline policy evaluation indicate that few-shot prompting, domain tuning, and grounded decision layers provide a viable path to deployable next-best-offer agents integrated into enterprise workflows.sap+3

References

[1]. P. Lewis et al., Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, NeurIPS, 2020.sendbird+1
[2]. T. B. Brown et al., Language Models are Few-Shot Learners, NeurIPS, 2020.knowmax+1
[3]. HarperValleyBank: A Domain-Specific Spoken Dialog Corpus, arXiv:2010.13929, 2020.advansappz+1
[4]. LinUCB and contextual bandits for personalized selection, overview and applications, True Theta, 2024.iopex+1
[5]. Deep Neural Network with LinUCB: A Contextual Bandit Approach, ACM, 2023.acm+1
[6]. Uplift Modeling: Making Predictive Models Actionable, Elder Research Whitepaper.sap+1
[7]. Response transformation and profit decomposition for revenue uplift, EJOR, 2020.marinsoftware+1
[8]. Enhancing Uplift Modeling in Multi-Treatment Marketing Campaigns, arXiv, 2024.findingtheta+1
[9]. What Is Retrieval-Augmented Generation (RAG), NVIDIA, 2025.mosaicx+1
[10]. How to Fine-Tune Zephyr-7B for Support Call Analysis, Predibase, 2024.forumvc+1
[11]. GDPR Article 22: Automated individual decision-making and profiling, gdpr-text.com and ICO guidance.ico+2
[12]. Automated Decision-Making Under the GDPR: Practical Cases from Courts and DPAs, FPF Report.fpf+1
[13]. India's Digital Personal Data Protection Act: Consent and rights overview, DLA Piper and Carnegie analyses.capgemini+2
[14]. Agentic AI in Telecommunications: Benefits & Use Cases, Salesforce, 2025.ey+1
[15]. How agentic AI is transforming IT: A CIO's guide, SAP, 2025.refuel+1
[16]. Agentic AI in Customer Service: Real-world applications, iOPEX, 2025.acm+1
[17]. Speech datasets and HarperValleyBank links, Stanford CS224S resources.startek+1
[18]. The State of AI: Global Survey, McKinsey, 2025.mckinsey
[19]. Analyzing Call Transcripts with LLMs: Topic extraction practices, Refuel.ai, 2024.refuel+1
[20]. Unlocking Voice of Customer from Call Transcripts with LLMs, Marin Software, 025.pub.towardsai+1
[21]. K. Kambhampati, "Explainable AI (XAI) for Trustworthy Cloud Security Decisions," Journal of Computational Analysis and Applications, vol. 31, no. 3, pp. 1605–1609, 2023, doi: 10.48047/jocaaa.2023.31.03.46.
[22]. K. Kambhampati, "AI-Driven Cloud User Validation for Secure Resource Allocation," International Journal of Intelligent Systems and Applications in Engineering, submitted May 9, 2022, revised Dec. 12, 2022, accepted Dec. 20, 2022.
[23]. K. Kambhampati, "The Sales Process and Service Process in Salesforce Using Agentic AI," Global Journal of Engineering and Technology, 2025.
[24]. K. Kambhampati, "Agentforce: Transforming CRM with Autonomous AI Agents in Salesforce," International Journal on Science and Technology, 2025.
[25]. K. Kambhampati, "Query-Efficient Multi-Agent Communication with Value-of-Information Policies and Typed Plan Messages," International Journal of Engineering, Management and Humanities (IJEMH), 2025.