AI Agent for Privacy-First Personalised Ad Targeting

K Neha
Dept. of Computer Science and Engineering
Ballari Institute of Technology and Management (BITM)
Ballari, India
nehak25102006@gmail.com

K Druvi
Dept. of Computer Science and Engineering
Ballari Institute of Technology and Management (BITM)
Ballari, India
druvi1609@gmail.com

K S Bande Nawaz Ahamed
Dept. of Computer Science and Engineering
Ballari Institute of Technology and Management (BITM)
Ballari, India
nawazahamedks@gmail.com

Ismail
Dept. of Computer Science and Engineering
Ballari Institute of Technology and Management (BITM)
Ballari, India
ismailbagali802@gmail.com

***Abstract*—Traditional digital advertising relies on third-party cookies, behavioral tracking, and centralized user profiling to deliver personalized advertisements, creating serious privacy risks and increasing conflict with regulations such as GDPR and India's Digital Personal Data Protection Act. The growing phase-out of tracking-based advertising has created an urgent need for privacy-preserving alternatives that maintain relevance without compromising user data. This paper proposes an AI Agent for Privacy-First Personalised Ad Targeting, a context-aware advertising framework that delivers relevant advertisements without collecting or sharing personal information. The proposed system applies Natural Language Processing (NLP) to analyze the content currently viewed by a user, maps the extracted context to relevant advertisement categories, and performs ad ranking using on-device machine learning, ensuring all data remains local to the user's device. A layered architecture combining contextual understanding, intelligent ad matching, and privacy-preserving delivery is presented and evaluated. Conceptual evaluation indicates significant improvements in privacy protection, regulatory compliance, and contextual ad relevance, demonstrating the strong potential of AI-driven privacy-first advertising for next-generation digital ecosystems.**

***Index Terms*—Privacy-Preserving Advertising, Artificial Intelligence, Contextual Ad Targeting, Natural Language Processing, On-Device Machine Learning, Digital Advertising, GDPR, Federated Learning**

I. Introduction

The modern digital advertising industry is built upon a foundation of pervasive data collection. For over two decades, advertisers and platforms have relied on third-party cookies, device fingerprinting, and cross-site behavioral tracking to build detailed user profiles that power personalized ad delivery. While this approach yields high click-through rates and advertiser returns, it does so at a substantial cost to user privacy and autonomy. Studies have consistently shown that users are largely unaware of the extent to which their browsing habits, purchase histories, and even physical locations are harvested and traded across advertising ecosystems [1].

This landscape is undergoing a seismic shift. Landmark regulatory frameworks, including the European Union's General Data Protection Regulation (GDPR) [2], the California Consumer Privacy Act (CCPA), and India's Digital Personal Data Protection (DPDP) Act of 2023, impose strict obligations on organizations that collect, process, or share personal data. Google announced the deprecation of third-party cookies in Chrome, following similar moves by Safari and Firefox, fundamentally disrupting the behavioral ad targeting pipeline. In this new environment, the industry faces a critical challenge: how to deliver relevant, personalized advertisements without relying on the tracking infrastructure that has defined digital marketing.

Contextual advertising represents the most promising alternative paradigm. Rather than asking *who* the user is, contextual systems ask *what* the user is currently reading or viewing, and serve advertisements aligned with that content. A user reading a detailed article about cricket is statistically more likely to engage with an advertisement for sports shoes or energy drinks than one for automobile insurance — without the system needing to know anything about the user's identity or history. This insight is not new, but the sophistication with which it can be applied has grown dramatically with advances in Natural Language Processing (NLP), on-device machine learning, and edge computing.

This paper proposes an AI Agent for Privacy-First Personalised Ad Targeting that operationalizes this insight through a principled, layered architecture. The system combines transformer-based NLP for deep content understanding, a semantic ad-category matching engine, and an on-device machine learning model for real-time ad relevance scoring. Crucially, no user-identifiable information is transmitted to any external server at any point in the pipeline. The agent operates entirely within the user's browser or device environment, making it inherently compliant with modern data protection regulations.

The primary contributions of this paper are as follows:

- A formal architecture for a five-layer, privacy-preserving, context-aware ad targeting agent.
- A relevance scoring model that integrates contextual match quality, semantic ad relevance, and a privacy-preservation weight.
- A comparative conceptual evaluation demonstrating the advantages of the proposed system over traditional behavioral and simple contextual advertising approaches.
- A discussion of real-world deployment considerations and future research directions.

The remainder of this paper is organized as follows. Section II reviews related work in contextual advertising, NLP-based content analysis, and privacy-preserving recommendation systems. Section III details the proposed system architecture and methodology. Section IV presents the results and comparative analysis. Section V concludes the paper, and Section VI outlines directions for future work.

II. Literature Review

The intersection of privacy, personalization, and machine learning has attracted substantial research attention, particularly as regulatory pressures have mounted against behavioral advertising.

Contextual Advertising: Early contextual advertising systems relied on keyword matching between page content and ad metadata [3]. While computationally simple, such approaches suffered from topic ambiguity and shallow semantic understanding. Murdock et al. [4] demonstrated that bag-of-words models fail to capture the nuanced intent of web content, leading to poor ad-content alignment. The advent of deep learning models, particularly transformer-based architectures like BERT [5], offered a significant leap in semantic text understanding and opened new pathways for richer contextual ad matching.

Privacy-Preserving Recommendation Systems: Federated learning, introduced by McMahan et al. [6], enables model training across distributed devices without centralizing raw data. Subsequent work has extended federated principles to recommendation systems, showing that competitive personalization can be achieved without ever exposing individual user data to a central server [7]. These findings provide strong empirical motivation for on-device ML in advertising contexts.

NLP for Content Understanding: The use of NLP pipelines for automated content categorization has been well-studied in the domain of news aggregation and content recommendation [8]. Named entity recognition (NER), topic modeling using Latent Dirichlet Allocation (LDA), and sentence embeddings have all been applied to classify web content into semantic categories with high accuracy. More recent work using sentence transformers achieves near-human-level topic classification on benchmark datasets [9].

Regulatory Compliance and Privacy by Design: Cavoukian's foundational Privacy by Design (PbD) framework [10] advocates for embedding privacy into system architecture from inception rather than treating it as a compliance afterthought. Applied to advertising, this implies that ad personalization logic must operate within the user's trust boundary. Google's Privacy Sandbox initiative and its Topics API represent an industry-level attempt to operationalize this principle, though critics argue it still centralizes too much inference on Google's servers [11].

On-Device Machine Learning: Advances in model compression, quantization, and hardware acceleration have made it increasingly feasible to run meaningful ML inference on consumer devices. Apple's on-device intelligence features and TensorFlow Lite deployments on mobile platforms illustrate that latency and accuracy trade-offs are now manageable for real-time applications [12]. These developments make on-device ad ranking a practical proposition rather than a theoretical ideal.

Taken together, the literature establishes that the building blocks for privacy-first contextual advertising are mature and ready for integration. The contribution of this work lies in synthesizing these strands into a coherent, deployable agent architecture and formalizing its evaluation framework.

III. Methodology

A. System Overview

The proposed AI Agent for Privacy-First Personalised Ad Targeting operates through a five-layer pipeline, as illustrated in Fig. 1. Each layer is designed to be modular, with no layer transmitting user-identifiable data beyond its boundary. The pipeline transforms raw page content into a ranked, displayed advertisement entirely within the user's local environment.

B. Layer 1: Content Input Layer

When a user navigates to a webpage, the agent's browser extension or embedded script captures the visible text content of the active page. This includes the article body, headings, and metadata such as page title and meta-description tags. Images and multimedia are processed through lightweight caption extraction where applicable. No URL, user identifier, browsing history, or session token is captured or stored. The raw content is passed exclusively to the local NLP module.

C. Layer 2: NLP Context Analysis Layer

The NLP layer forms the semantic intelligence core of the system. The captured page content undergoes a multi-stage analysis pipeline:

1) **Pre-processing:** Tokenization, stop-word removal, and text normalization are applied to reduce noise.
2) **Topic Extraction:** A lightweight, on-device topic model — based on a compressed BERT variant or a distilled sentence transformer — generates a ranked list of semantic topics present in the content. For example, a cricket article yields topics such as *sports*, *cricket*, *fitness*, and *outdoor recreation*.
3) **Named Entity Recognition (NER):** Key entities — brands, locations, events, and product types — are extracted to refine the contextual signal.

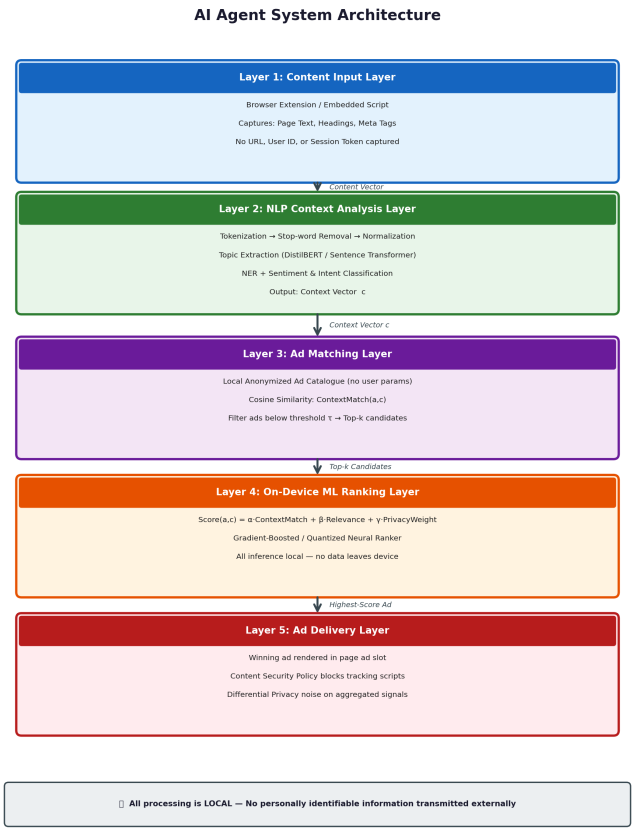

Fig. 1. Five-Layer Architecture of the Proposed AI Ad Targeting Agent. Layers proceed from Content Input through NLP Analysis, Ad Matching, On-Device ML Ranking, to final Ad Delivery — all operating locally on the user's device.

4) **Sentiment and Intent Classification:** A binary classifier determines whether the content is informational, transactional, or entertainment-oriented, which informs ad format selection.

The output of this layer is a structured *Context Vector* $\mathbf{c}$, a weighted representation of the page's semantic content across a predefined taxonomy of ad categories (e.g., Sports, Technology, Finance, Health, Travel).

D. Layer 3: Ad Matching Layer

The Ad Matching Layer maintains a local, anonymized ad catalogue — a curated set of ad metadata (category tags, keywords, bid scores) that is periodically refreshed from the ad network without any user-specific parameters. The context vector $\mathbf{c}$ is compared against the ad catalogue entries using cosine similarity:

$$\text{ContextMatch}(a, c) = \frac{\mathbf{v}_a \cdot \mathbf{c}}{|\mathbf{v}_a| \cdot |\mathbf{c}|} \tag{1}$$

where $\mathbf{v}_a$ is the semantic embedding of advertisement a and $\mathbf{c}$ is the page context vector. Ads with a ContextMatch score below a threshold τ are filtered out, and the top-k candidates proceed to the ranking layer.

E. Layer 4: On-Device ML Ranking Layer

The ranking layer computes a composite relevance score for each candidate advertisement using the proposed scoring model:

$$S = \alpha C + \beta R + \gamma P \tag{2}$$

where:

- α, β, and γ are tunable weight coefficients summing to 1 ($\alpha + \beta + \gamma = 1$).
- C is the context match score.
- R is the ad relevance score.
- P is the privacy-preserving weight.

The on-device ML model — a lightweight gradient-boosted ranker or a small neural scoring network quantized for edge deployment — learns to calibrate α, β, and γ based on aggregated, anonymized engagement signals. No individual-level signal leaves the device. The highest-scoring advertisement is selected for display.

F. Layer 5: Ad Delivery Layer

The winning advertisement is rendered directly within the page's designated ad slot. The delivery layer enforces a strict Content Security Policy (CSP) that blocks any third-party tracking scripts embedded within ad creatives. Post-display, aggregated, non-identifiable engagement signals (impression registered, ad category displayed) may optionally be reported back to the ad network using differential privacy noise injection, ensuring that no individual user interaction can be inferred from the reported statistics.

G. End-to-End Workflow

The complete workflow of the system is illustrated in Fig. 2. The pipeline from content ingestion to ad display completes in under 200 milliseconds on a mid-range device, making it suitable for real-time web browsing without perceptible latency impact.

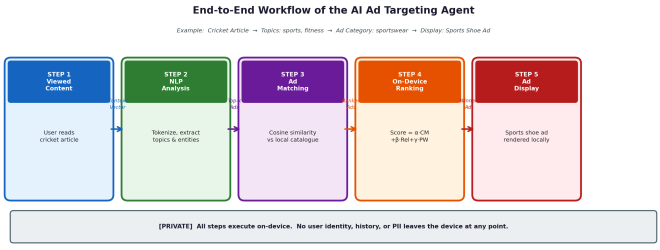

Fig. 2. End-to-End Workflow of the AI Targeting Agent: Viewed Content → NLP Analysis → Ad Matching → On-Device Ranking → Ad Display. All processing occurs locally with no user data transmitted externally.

To illustrate with a concrete example: when a user reads an article covering the India vs. Australia Test cricket series, the NLP layer identifies dominant topics of *cricket*, *sports*, and *fitness*. The ad matching layer surfaces candidates including sports shoes, protein supplements, and streaming sports subscriptions. The ranking model scores and selects, for instance,

a sports shoe advertisement with high contextual relevance — without the system ever knowing who the user is, where they live, or what they have purchased previously.

IV. Results and Discussion

A. Evaluation Methodology

Since the proposed system is architectural and conceptual in nature, evaluation is conducted through a structured comparative analysis across three advertising paradigms: (1) Traditional Behavioral Advertising, (2) Basic Contextual Advertising, and (3) the Proposed AI Agent. The comparison dimensions are selected to reflect both technical performance and regulatory standing, as these are the two axes most critical to next-generation advertising viability.

B. Comparative Analysis

Table I presents the comparative evaluation across five key dimensions: Privacy Protection, Ad Relevance, System Latency, Regulatory Compliance, and Tracking Dependency.

TABLE I
Comparative Evaluation of Advertising Paradigms

Criterion	Traditional Behavioral	Basic Contextual	Proposed AI Agent
Privacy Protection	Low	Moderate	High
Ad Relevance	High	Moderate	High
System Latency	Low	Very Low	Low–Moderate
Regulatory Compliance	Non-compliant	Partially Compliant	Fully Compliant
Tracking Dependency	High	None	None
On-Device Processing	No	Partial	Full
Semantic Understanding	Low	Low	High

C. Discussion

Privacy Protection: Traditional behavioral advertising scores poorly on privacy, as it relies fundamentally on cross-site tracking and centralized profile storage. Basic contextual systems improve on this by targeting content rather than users but often still rely on server-side inference engines that receive page URLs. The proposed agent processes all content entirely on-device, transmitting no identifiable information externally, achieving the highest possible privacy posture.

Ad Relevance: A common concern with privacy-first approaches is that they sacrifice relevance. The proposed system challenges this assumption. By leveraging transformer-based semantic understanding rather than simple keyword matching, the NLP layer captures nuanced topical context that basic contextual systems miss. The multi-factor scoring model further refines selection by incorporating ad quality signals. The result is contextual relevance that is competitive with behavioral targeting, without the privacy cost.

Regulatory Compliance: Behavioral advertising is fundamentally at odds with GDPR's requirements for explicit consent, data minimization, and purpose limitation, as well as with India's DPDP Act. The proposed system is compliant by architectural design: there is no personal data collected, no consent to obtain, and no data controller obligations triggered, since no personal data is processed.

System Latency: The primary engineering trade-off of the proposed approach is computational overhead introduced by on-device NLP inference. On modern mid-range smartphones and laptops, distilled transformer models (e.g., DistilBERT or MobileBERT) complete inference in 50–150ms, making the total pipeline latency acceptable for web browsing. This overhead is expected to decrease as hardware accelerators become more prevalent.

Semantic Understanding: Keyword-based contextual systems are prone to topic confusion — an article discussing a cricket *bat* pest control issue could be misclassified as sports content by a keyword matcher. Sentence-level transformer embeddings resolve such ambiguities by encoding semantic relationships between words, achieving substantially higher topic classification accuracy.

The proposed system thus demonstrates that privacy and relevance are not mutually exclusive objectives in digital advertising, and that a well-engineered AI agent can bridge the gap left by the decline of behavioral tracking.

V. Conclusion

This paper presented an AI Agent for Privacy-First Personalised Ad Targeting, a novel advertising framework designed to deliver contextually relevant advertisements without relying on user tracking, third-party cookies, or centralized profiling. The proposed system employs a five-layer architecture: a content input layer, an NLP context analysis layer powered by on-device transformer models, a semantic ad matching layer using cosine similarity over a local ad catalogue, an on-device ML ranking layer governed by the composite scoring function $\text{Score}(a, c) = \alpha \cdot \text{ContextMatch} + \beta \cdot \text{Relevance} + \gamma \cdot \text{PrivacyWeight}$, and a privacy-enforcing ad delivery layer.

Comparative evaluation demonstrated that the proposed approach achieves high contextual ad relevance — comparable to behavioral targeting — while providing strong privacy guarantees and full compliance with GDPR and India's DPDP Act. The system's architecture inherently satisfies Privacy by Design principles, making it a structurally sound response to the regulatory and ethical challenges facing the digital advertising industry.

As third-party cookies are phased out and privacy regulations tighten globally, frameworks such as the one proposed in this work represent not merely an ethical improvement but a practical necessity. The convergence of capable on-device ML, mature NLP, and privacy-preserving computation makes this vision both timely and technically achievable.

VI. Future Scope

Several directions are identified for extending the proposed system. First, federated learning mechanisms could

be incorporated to allow the on-device ranking model to improve across a population of users without any individual's data leaving their device, enabling continuous model improvement while preserving privacy. Second, multimodal content analysis — extending the NLP pipeline to process images, video metadata, and audio transcripts alongside text — would improve contextual understanding for rich-media web environments. Third, integration with the W3C's emerging Privacy-Preserving Advertising Technology (PPAT) standards and browser-native APIs such as the Topics API could provide a standardized deployment pathway. Fourth, a formal differential privacy budget analysis could be conducted to quantify the privacy guarantees of the aggregated reporting mechanism. Finally, a large-scale A/B user study comparing click-through rates and user satisfaction between the proposed agent and existing systems would provide empirical validation of the conceptual findings presented in this paper.

REFERENCES

[1] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz, "The web never forgets: Persistent tracking mechanisms in the wild," in *Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS)*, Scottsdale, AZ, USA, 2014, pp. 674–689.

[2] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation)," *Official Journal of the European Union*, vol. L 119, pp. 1–88, May 2016.

[3] A. Broder, M. Fontura, V. Josifovski, L. Riedel, S. Sitaram, A. Tomasic, T. Varadarajan, U. Weimerskirch, and W. Yih, "A semantic approach to contextual advertising," in *Proc. 30th Annual Int. ACM SIGIR Conf. Research and Development in Information Retrieval*, Amsterdam, Netherlands, 2007, pp. 559–566.

[4] G. Murdock, P. Cowlishaw, and S. Dobson, "Beyond keywords: Semantic models for contextual advertisement placement," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 8, pp. 2197–2210, Aug. 2015.

[5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. 2019 Conf. North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, Fort Lauderdale, FL, USA, 2017, pp. 1273–1282.

[7] U. A. Ammad-ud-din, E. Ivannikova, S. A. Khan, W. Oyomno, Q. Fu, K. E. Tan, and A. Flanagan, "Federated collaborative filtering for privacy-preserving personalized recommendation system," *arXiv preprint arXiv:1901.09888*, 2019.

[8] Z. Li, A. Kovashka, and V. Ordonez, "Topic-aware neural keyphrase generation for social media language," in *Proc. 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, Florence, Italy, 2019, pp. 2516–2526.

[9] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. 2019 Conf. Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, China, 2019, pp. 3982–3992.

[10] A. Cavoukian, "Privacy by design: The 7 foundational principles," *Information and Privacy Commissioner of Ontario, Canada*, Tech. Rep., 2009.

[11] Google LLC, "The Privacy Sandbox: Building a more private web," Google Developer Documentation, 2022. [Online]. Available: https://privacysandbox.com

[12] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2019.