Understanding How Recommender Systems Shape Consumer Behavior Using Implicit Feedback

Aditya Dev

Daly College Indore, M.P., India

Abstract: ***This paper explores how major technology firms like Google, Amazon, Meta, and Microsoft apply artificial intelligence and machine learning to understand and influence consumer behavior using implicit feedback such as clicks and watch time. By examining collaborative filtering, content-based, hybrid, and neural network–driven recommendation systems, it highlights both their technical foundation and real-world scalability. Through detailed case studies, the study analyzes the implications of personalized digital experiences and their effect on consumer decision- making. It also critiques the ethical dimensions of data usage, addressing concerns around privacy, autonomy, and algorithmic bias. Ultimately, the paper argues that while recommender systems boost personalization and commercial success, their sustainable future hinges on responsible, transparent deployment.***

Keywords: Consumer behavior modeling, recommender systems, data privacy, deep learning, personalization

1. Introduction

Major technology companies deploy a wide range of AI-driven products that shape how users interact with digital platforms. Examples include Meta's AI Sandbox, Google Ads platform and Search engine, Amazon's product recommendation engine, and Netflix's personalized content recommender. Despite their different purposes, all of these products rely on *recommender systems*—algorithms that analyze user behavior and preferences to recommend relevant content, including products and services. Their influence is significant, determining much of what users see online and subtly shaping decision-making. In a study by Walters and Jamil, 39% of items in a typical grocery shopping basket were special offers, and 30% of surveyed customers reported being highly influenced by coupons and promotions [1]. Such evidence highlights the persuasive power of recommendation mechanisms. Understanding how these systems work benefits both consumers—by promoting transparency—and FMCG businesses—by enabling data-driven strategies.

For the context of this paper, we will define a consumer as any user interacting with a digital platform whose behavior, preferences, or actions can be tracked and modeled to deliver personalized content, recommendations, or advertisements using machine learning. Their consumption is their interaction with digital content that is directed towards them, be it viewing, clicking, scrolling, liking, watching, or buying. Their effect on consumer behavior is pivotal in understanding the influence of these big tech companies. The reasons are as follows [2]:

1) **Personalized Shopping Experiences**: Real-time recommendations can be made to provide a personalized shopping experience, which can encourage them to purchase and hence improve the overall profit of these organizations.
2) **Strategic Business Decision-Making**: Such insights can be useful to e-commerce organizations to assess product popularity and monitor user activity.
3) **Better Product Development and Targeting**: These studies help the organizations in knowing their customers and incorporating targeted marketing techniques to increase their customer base and profits.
4) **Awareness of Influence**: Big tech platforms like Amazon and Netflix use algorithms to track our behavior and suggest products. Knowing this helps people become more aware of how their choices are shaped, often without realizing it.
5) **Data Ownership and Privacy**: Learning how companies collect and use personal data for predictions helps people value their privacy. It encourages more mindful sharing of personal information online.

2. Technical Background

A. Technologies Used in Consumer Analytics

Big technology companies (such as Amazon, Google, Meta, Netflix, and Apple) leverage a wide range of advanced technologies to understand consumer behavior and deliver personalized experiences. The key technologies used include machine learning, recommender systems, natural language processing, and deep learning architectures. These systems help companies optimize digital products, targeted advertising, and recommendation engines.

1) *Machine Learning:* Machine learning (ML) is a subset of artificial intelligence that allows systems to learn from data without being explicitly programmed. It is a foundational technology behind many recommendation and personalization algorithms.

- **Supervised Learning:** This involves training a model on labeled data. Big tech companies use supervised learning to predict user preferences, click-through rates, or product ratings. (i) Regression is a predictive modeling technique that examines the relationship between a dependent variable and an independent variable. (ii) Classification is the technique in which the algorithm learns from the data input given to it and then uses this learning to classify and produce new observation. [3]

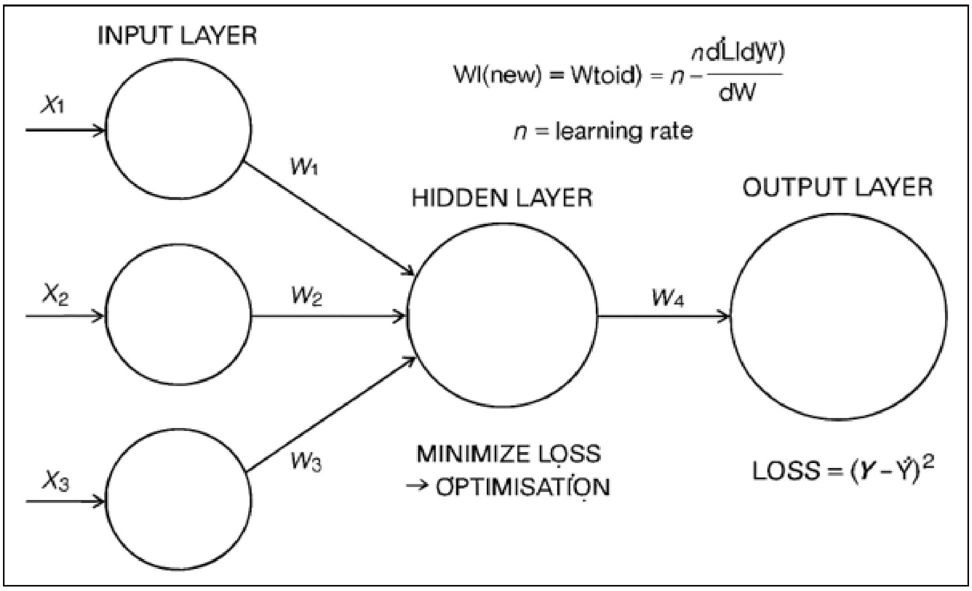

Figure 1: A schematic diagram of an artificial neural network showing the input, hidden, and output layers with weight updates.

- **Unsupervised Learning:** Used to find hidden patterns or groupings in data, such as customer segmentation based on purchasing behavior.
- **Reinforcement Learning:** In this method, agents learn optimal actions through rewards and penalties. It is often used in real-time recommendation systems and personalized content delivery.
- **Semi-supervised and Self-supervised Learning:** These approaches leverage partially labeled or unlabeled data, which is beneficial given the massive volume of user data big tech platforms manage.

2) *Recommender Systems:* Recommender systems are a specific application of machine learning that suggest products, media, or services to users based on various signals.

- **Collaborative Filtering:** This method recommends items by identifying similarities between users or items. For example, Amazon uses item-based collaborative filtering to suggest products that similar users have purchased.
- **Content-Based Filtering:** Recommends items similar to what a user has liked in the past, based on product features or descriptions.
- **Hybrid Models:** Combine collaborative and content-based approaches to improve accuracy and robustness.

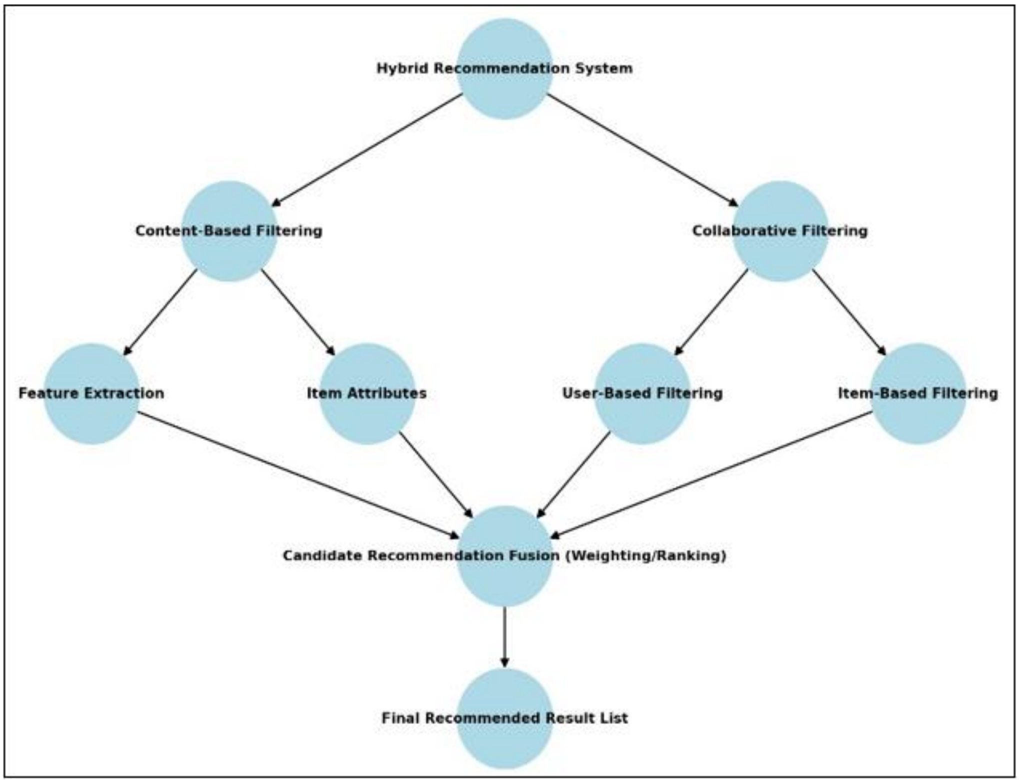

Figure 2: Architecture of a Hybrid Recommendation System

Paper ID: SR25909210040 DOI: https://dx.doi.org/10.21275/SR25909210040

- **Deep Learning-based Systems:** Modern recommender systems also use neural networks to capture complex relationships in user behavior and preferences.

3) *Other Supporting Technologies:*

- **A/B Testing and Feedback Loops:** Companies continuously run experiments to refine algorithms, interfaces, and marketing strategies based on real-time user feedback.
- **Clustering and Segmentation:** These help segment users based on their behavioral patterns. This enables targeted marketing and personalized recommendations.
- **Big Data Infrastructure:** Technologies like Hadoop, Spark, and cloud platforms enable real-time processing and storage of massive datasets generated by millions of users.

These tools allow tech companies to analyze and predict consumer behavior more effectively than ever before. [4]

3. Big Tech Review

1) Meta

Meta Platforms Inc. runs targeted advertising across Facebook, Instagram, Messenger, and WhatsApp, using machine learning to optimize reach and relevance. Its diverse ad formats, dual campaign tools (Ads Manager and Business Suite), and automation features like Advantage+ enable flexible, performance-driven marketing for businesses of all sizes.

Meta Platforms Inc. employs advanced AI-driven technologies to enhance advertising performance and user engagement across its platforms. Below are three key technologies that demonstrate its approach:

a) *Advantage and Advantage+:* These are automated optimization tools for ad campaigns that leverage machine learning to streamline the advertising process. They automatically adjust targeting, bidding, and ad placements to maximize campaign performance, giving advertisers a powerful, hands-off optimization capability. [5]

   From the consumer perspective, Advantage and Advantage+ ensure that ads reach users at the right moment and in the right format, tailoring exposure to those most likely to respond. This level of personalization influences user decision-making by reinforcing purchasing intent, increasing the likelihood of impulse buying, and making ad experiences feel more relevant and less intrusive.

b) *Andromeda:* This proprietary machine learning system is designed for ad recommendation retrieval. It uses deep neural networks combined with hierarchical indexing to improve both recall and the quality of ad recommen- dations. Reported results include a 6% increase in recall, an 8% improvement in ad quality, and a maximized return on investment for advertisers. [6]

   For consumers, Andromeda's improvements in recall and ad quality translate into more accurate and compelling ad exposures. Instead of being shown irrelevant or generic promotions, users consistently encounter ads that align with their interests and browsing patterns, thereby subtly shaping preferences and guiding their future purchase behaviors.

c) *AdLlama:* This large language model (LLM) is specialized in generating ad text variations. Using reinforce- ment learning with performance feedback (RLPF), it produces optimized and diverse ad content. Its deployment resulted in a 6.7% increase in advertiser-level click-through rates (CTR) and an 18.5% increase in the number of ad variations created. [7]

   From the consumer's viewpoint, AdLlama increases engagement by presenting ad content in diverse tones and styles that better resonate with different audiences. This not only sustains user attention but also fosters a perception of novelty, making consumers more likely to engage with products and services they might have otherwise overlooked.

d) *AI Sandbox:* This generative AI tool, called Text Variation, helps advertisers rapidly create multiple ad copy variations from a single prompt. By generating alternative phrasings, styles, and tones, it streamlines A/B testing and enables message tailoring for specific audiences, thereby improving campaign performance [8], [9].

   Disclaimer: The example (Listing 1) is solely for illustrative purposes and adapts Hugging Face's open-source `transformers` *library [10], [11].*

```python
from transformers import pipeline

# Load a text generation model
generator = pipeline("text-generation", model="gpt2")

# Original ad text
prompt = "Discover the best deals on eco-friendly home products."

# Generate multiple variations
variations = generator(prompt, num_return_sequences=3, max_length=30)

# Print each variation
for i, var in enumerate(variations, 1):
    print(f"Variation {i}: {var['generated_text']}")
```

Listing 1: Generating ad copy variations using Hugging Face's transformers library

In Listing 1, the `pipeline` function from the `transformers` library is used to load a pre-trained text generation model (`gpt2`). The variable `prompt` stores the original ad text to be rephrased. The `generator` object is then called with parameters specifying the input prompt, the number of alternative sequences to return (`num_return_sequences=3`), and the maximum length for each generated text (`max_length=30`). The resulting list of `variations` contains different phrasings of the original ad text. Finally, a loop iterates through these variations and prints each one in a numbered format.

In practice, Meta's production system goes beyond this simple example by incorporating reinforcement learning with performance feedback (RLPF) to rank and prioritize ad variations that have historically achieved higher engagement or conversion rates. This helps generate diverse content optimized for measurable advertising goals.

For consumers, AI Sandbox means exposure to ad messages that feel more personalized and contextually relevant. By continuously testing and refining ad copy variations, it increases the likelihood that users will find ads appealing and trustworthy, thereby shaping browsing behavior and subtly nudging them toward purchase decisions.

2) Google

Google has reported continuous growth in advertising revenue, driven in part by its application of machine learning and artificial intelligence technologies. In its 2024 Annual Report, Google highlighted significant investments in AI to improve ad personalization, which increases the relevance of ads served to users and enhances marketing performance for advertisers [12].

a) *Deep Learning for User Behavior Analysis:* Google employs deep learning to analyze user interactions across platforms such as Search, YouTube, and Maps. These models infer intent and preferences in real-time, enabling the system to recommend content and products more accurately [12].
   From the consumer perspective, this means that searches, video recommendations, and even map suggestions feel increasingly personalized. By learning from browsing histories and engagement patterns, Google nudges users toward products, videos, or services they are statistically more likely to click on, reinforcing consumption habits and sustaining platform engagement.
b) *Reinforcement Learning in Smart Bidding:* Google Ads integrates reinforcement learning in Smart Bidding, dynamically adjusting ad strategies based on contextual signals, conversions, and user interactions [12]. This allows advertisers to optimize campaigns in real time for outcomes such as clicks or conversions.
   For consumers, Smart Bidding ensures that the ads they see are not random but tailored to their browsing context and purchasing likelihood. Over time, this increases the relevance of sponsored content in Search and YouTube, subtly influencing consumer purchasing decisions and encouraging faster, more impulsive conversions.

3) Amazon

Amazon leverages AI techniques to understand consumer behavior and deliver highly personalized shopping experiences. Two key methods are used: sentiment analysis of product reviews to extract opinions, and recommender systems to suggest items aligned with user preferences [2].

a) *User-based Collaborative Filtering:* This method identifies customers whose purchase and rating histories overlap with the target user. By calculating similarity scores, it recommends products that similar users have engaged with but the current user has not [2].
   From the consumer perspective, this creates a sense of social validation—users feel reassured when shown items that similar customers have purchased. This not only increases trust in recommendations but also encourages users to try new products they might not have discovered independently.
b) *Item-based Collaborative Filtering:* Instead of focusing on users, this method analyzes item-to-item similarities. It recommends products similar to those a consumer has already interacted with positively, a design that improves scalability for Amazon's massive catalog [2].
   For consumers, this results in smoother browsing and a "continuous discovery" effect—where one purchase leads to further suggestions closely tied to their preferences. Over time, this deepens consumer engagement with the platform and subtly drives repeat buying behavior.

4) Microsoft

Microsoft employs advanced AI technologies to improve the scalability, efficiency, and accuracy of recom- mender systems. Two key approaches include neural network–based recommendation models for capturing complex user–item relationships, and distributed optimization techniques that enable the training of extremely large models across multiple GPUs. Together, these methods allow Microsoft to deliver highly personalized recommendations while handling the computational demands of modern AI systems.

Neural Collaborative Filtering: In Microsoft's open-source *Recommenders* repository, Neural Collaborative Filtering (NCF) is implemented as a deep learning approach to recommender systems. NCF replaces the inner product in traditional matrix factorization with a multi-layer perceptron, allowing it to capture complex, non-linear user–item interactions. The following Python snippet (Listing 2) from the repository demonstrates how to train and evaluate an NCF model [13].

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

```python
from reco_utils.recommender.ncf.ncf_singlenode import NCF
from reco_utils.dataset import movielens

# Load dataset
data = movielens.load_pandas_df(size="100k")
train, test = movielens.split_train_test(data, ratio=0.75)

# Initialize and train NCF model
model = NCF(
    user_col='userID',
    item_col='movieID',
    rating_col='rating',
    seed=42,
    epochs=10,
    batch_size=256,
    learning_rate=0.001
)
model.fit(train)

# Evaluate model
metrics = model.evaluate(test, top_k=10)
print(metrics)
```

Listing 2: Training and evaluating an NCF model from Microsoft's Recommenders repository

In Listing 2, the `movielens.load_pandas_df` function loads the dataset, which is then split into training (75%) and testing (25%). An `NCF` model is created with specified parameters and trained on the data before being evaluated using top-10 recommendation metrics. The printed results reflect its accuracy and personalization capabilities.

For consumers, this results in more relevant and timely recommendations on Microsoft-powered platforms, reducing time spent searching and increasing the likelihood of discovering content or products that match their interests.

1) Zero Redundancy Optimizer: Microsoft's *DeepSpeed* library includes the Zero Redundancy Optimizer (ZeRO), which enables the training of extremely large models by partitioning model states—including parameters, gradients, and optimizer states—across multiple GPUs. This reduces memory usage and allows models with hundreds of billions of parameters to be trained efficiently [14].

```python
import deepspeed
import torch.nn as nn

# partitioning parameters across the data parallel group.
with deepspeed.zero.Init(
    data_parallel_group=mpu.get_data_parallel_group(),
    remote_device=get_args().remote_device,
    enabled=get_args().zero_stage==3
):
    # Activates ZeRO Stage 3
    model = GPT2Model(num_tokentypes=0, parallel_output=True)
```

Listing 3: Initializing a model with ZeRO Stage 3

```python
import torch

self.position_embeddings = torch.nn.Embedding(...)
# This context manager ensures the embedding weights are gathered for initialization and then correctly
    re-partitioned by DeepSpeed.
with deepspeed.zero.GatheredParameters(
    self.position_embeddings.weight,
    modifier_rank=0):
    # This initialized the position embeddings using the chosen method.
    self.init_method(self.position_embeddings.weight)
```

Listing 4: Managing embedding initialization with ZeRO

In Listings 3 and 4, `deepspeed.zero.Init` partitions parameters across GPUs, while `GatheredParameters` ensures consistent initialization before redistributing embeddings. These optimizations minimize GPU memory usage and enable large-scale training.

For consumers, ZeRO's impact is indirect but significant: it supports the creation of more powerful recommender models and AI services that feel faster, smarter, and more adaptive. By making large-scale personalization feasible, Microsoft ensures that users consistently receive highly tailored recommendations and seamless digital experiences.

4. Analysis

Firstly, all of the companies use machine learning— often deep learning— to tailor recommendations or ads to individual users. This stems from a shared goal among big tech companies: to recognize and analyze consumer behavior in order to improve relevance and engagement.

Next, they all operate core recommendation engines that drive content discovery and make accurate predictions of what the user is most likely to interact with next.

Companies such as Google and Microsoft manage massive volumes of data, which necessitates rapid model training at scale. For this reason, they employ distributed training, described below.

A. Distributed Training

Distributed training trains machine learning models across multiple GPUs or servers using *data parallelism* (splitting data) and *model parallelism* (splitting parameters) for faster, scalable training without exceeding hardware limits.

Once distilled, these technical approaches form the backbone of the personalization strategies deployed by big tech companies. By processing and interpreting vast amounts of implicit feedback—such as clicks, watch time, and scrolling behavior—these systems can deliver highly targeted recommendations that directly influence consumer decision-making. While such precision can enhance user satisfaction and business performance, it also raises important questions about autonomy, manipulation, and privacy. These issues, along with their broader societal implications, are discussed in detail in the *Ethical Impact* section.

5. Future Direction

The leading technology firms—Google (Alphabet), Amazon, Meta, and Microsoft—are driving innovation through strategic investments in artificial intelligence, cloud infrastructure, and platform integration. These initiatives aim to improve scalability, personalization, and cross-platform experiences, positioning each company at the forefront of the evolving digital economy. For consumers, this results in faster services, smarter recommendations, and more immersive online environments. Table I provides a Q2 2025 performance overview, reflecting the scale of resources committed to shaping future technological directions.

These strategic shifts are reshaping consumer interaction with digital platforms. AI-powered personalization is enabling services that anticipate user needs, while advances in cloud computing are delivering faster, more reliable access to applications from anywhere. Immersive technologies such as virtual and augmented reality are expanding the scope of online experiences, from social interaction to shopping and professional collaboration. Looking ahead, as generative AI and automation advance, digital and physical boundaries will become increasingly blurred between physical and digital environments, creating more intuitive, data-driven, and interconnected consumer experiences.

Table I: Q2 2025 Financial Snapshot for Google, Meta, Microsoft, and Amazon

Metric	Google (Alphabet) (Q2 2025)	Meta (Q2 2025)	Microsoft (FY25 Q2, ended Dec 2024)	Amazon (Q2 2025)
Consolidated Revenue	\$96.4 billion	\$47.52 billion	\$69.6 billion	\$167.7 billion
Cloud Revenue	\$13.6 billion (Google Cloud)	N/A	\$25.5 billion (Intelligent Cloud)	\$30.9 billion (AWS)
Operating Income	\$31.271 billion	\$20.441 billion	\$31.7 billion	\$19.2 billion
Net Income	\$28.196 billion	\$18.337 billion	\$24.1 billion	\$18.2 billion
Capital Expenditures (Annual/Outlook)	~\$85 billion (2025 Outlook)	\$66–72 billion (2025 Outlook)	Significant, not fully disclosed (FY25)	\$103 billion (2025 Outlook)

6. Ethical Impact

The integration of AI-driven recommender systems into digital platforms has significant ethical implications. While these systems enhance personalization and user engagement, they also introduce a range of societal and individual risks. Key concerns include:

1) **Privacy and Surveillance:** Recommender systems rely heavily on large-scale collection and analysis of consumer data. This can create a perception of constant surveillance and raise concerns over how personal information is stored, shared, and repurposed [15].
2) **Loss of Autonomy:** Continuous tracking and prediction of user behavior can subtly manipulate consumer decisions, undermining individual autonomy and control over one's digital presence [16].
3) **Bias and Fairness:** Algorithmic bias in recommender systems can lead to unequal treatment of users, amplify existing societal biases, and contribute to digital divides. Such biases may emerge from skewed training data, reinforcement of popular content, or systemic underrepresentation of minority groups [17], [18].
4) **Lack of Human Oversight:** Many large-scale recommender systems operate with minimal human monitoring, relying almost entirely on automated decision-making. Without adequate oversight, errors, harmful content, or biased recommendations can persist unchecked [19].
5) **Erosion of Trust:** For some users, the possibility of their data being misused for targeted advertising or behavioral manipulation generates insecurity and distrust. This, over time, can foster public resistance and reduce the effectiveness of these systems, despite their commercial advantages for big technology companies [20].

Addressing these ethical challenges requires a combination of transparency, algorithmic auditing, bias mitigation strategies, and the integration of human oversight into automated systems.

7. Conclusion

Modern recommendation engines use artificial intelligence to offer you highly tailored experiences, helping you quickly find products, services, and content that you might like. However, they depend on gathering and analyzing a constant stream of your personal data. This practice can make some people feel uneasy, as if their privacy is being invaded.

The feeling of being constantly watched, along with the risk of data being misused for things like targeted ads, political manipulation, or security breaches, can create a sense of insecurity. For some, the comfort of knowing they control their personal information is more important than the convenience of these systems.

This situation raises critical ethical questions about whether people truly understand and consent to how their data is used. It's a debate about how to balance technological innovation with the fundamental right to individual privacy. While these systems are very profitable for big tech companies, their future success may hinge on how well they can handle these privacy issues

References

[1] N. Chabane, A. Bouaoune, R. Tighilt, M. Abdar, A. Boc, E. Lord, N. Tahiri, B. Mazoure, U. R. Acharya, and V. Makarenkov, "Intelligent personalized shopping recommendation using clustering and supervised machine learning algorithms," *PLOS ONE*, vol. 17, no. 12, p. e0278364, 2022. [Online]. Available: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0278364

[2] K. Meghana, K. N. V. Rao, and N. Geethanjali, "Consumer behavior analytics using machine learning algorithms," *International Journal of Research in Engineering, Science and Management*, vol. 5, no. 4, pp. 118–121, 2022. [Online]. Available: https://www.ijresm.com/Vol.5 2022/Issue 4/HTML/24.html

[3] B. Ramzan, I. S. Bajwa, N. Jamil, R. U. Amin, S. Ramzan, F. Mirza, and N. Sarwar, "An intelligent data analysis for recommendation systems using machine learning," *Scientific Programming*, vol. 2019, pp. Article ID 5 941 096, 20 pages, 2019.

[4] L. Sun, Q. Chen, J. Wang, W. Li, and J. Zhang, "Blockchain-based data security and privacy protection in internet of things: A comprehensive survey," *Journal of Computer Science and Systems Research*, vol. 4, no. 2, pp. 45–60, 2022. [Online]. Available: https://thescipub.com/abstract/jcssr.2022.45.60

[5] DataFeedWatch, "Facebook dynamic ads for retail: What you need to know," https://www.datafeedwatch.com/blog/facebook-dynamic-ads-for-retail-what-you-need-to-know, 2024, accessed: 2025-08-09.

[6] Meta Platforms Inc., "Andromeda and advantage+: Automation for the next generation of personalized ads retrieval," https://engineering.fb.com/2024/12/02/production-engineering/meta-andromeda-advantage-automation-next-gen-personalized-ads- retrieval-engine/, accessed: 2025-08-09.

[7] L. Sun *et al.*, "Improving generative ad text on facebook using reinforcement learning," *ResearchGate*, 2024, accessed: 2025-08-09.

[8] P. Adams, "Meta broadens access to generative ai tools, seeking to save marketers time," *Marketing Dive*, May 2023, "Meta's AI Sandbox . . . lets advertisers create multiple variations of text copy". [Online]. Available: https://www.marketingdive.com/news/Meta-launches-AI-sandbox-tool-social-media-advertising/650117/

[9] K. Hines, "Meta announces ai-powered tools to streamline ad processes," *Search Engine Journal*, May 2023, "Text Variation automatically creates multiple versions of an advertiser's copy". [Online]. Available: https://www.searchenginejournal.com/meta-ai-powered-ad-tools/ 486778/

[10] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer,

[11] P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. [Online]. Available: https://aclanthology.org/2020.emnlp-demos.6

[12] Hugging Face, "Transformers documentation," https://huggingface.co/docs/transformers, 2024, accessed: 2025-08-16.

[13] Google LLC, "Alphabet inc. 2024 annual report," Alphabet Inc., Tech. Rep., 2025. [Online]. Available: https://abc.xyz/investor/static/pdf/ 2024 Annual Report.pdf

[14] Microsoft, "Recommenders: Examples and best practices for building recommendation systems," https://github.com/microsoft/ recommenders, 2023, accessed: 2025-08-14.

[15] S. Rajbhandari, J. Rasley, O. Ruwase, and Y. He, "Zero: Memory optimizations toward training trillion parameter models," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2020, pp. 1–16.

[16] S. Milano, M. Taddeo, and L. Floridi, "Recommender systems and their ethical challenges," *AI & Society*, vol. 35, pp. 957–967, 2020.

[17] E. Carmi and S. Yates, "Digital surveillance, privacy and the pandemic: An analysis of the uk's nhs covid-19 app," *Social Media+ Society*, vol. 6, no. 4, 2020.

[18] M. K. Lee, K. Hosanagar, and H. S. Nair, "Algorithmic bias in recommender systems: Examining bias amplification and mitigation," in *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 2019.

[19] M. D. Ekstrand, M. Tian, I. Azpiazu, J. Ekstrand, D. McNeill, O. Anuyah, and M. S. Pera, "Exploring author gender in book rating and recommendation," in *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, 2018, pp. 224–232.

[20] S. Barocas and A. D. Selbst, "Big data's disparate

Paper ID: SR25909210040 DOI: https://dx.doi.org/10.21275/SR25909210040

impact," in *Data and Discrimination: Collected Essays*. MIT Press, 2016, pp. 93–128.

[21] A. Acquisti, L. Brandimarte, and G. Loewenstein, "Privacy and human behavior in the age of information," *Science*, vol. 347, no. 6221, pp. 509–514, 2015.