

Bias Resistant Retrieval Augmented Generation: A Clustering and BiQ Driven Approach with Equitable AI

Vignesh K^{1*}; Sharanjey G²; Pranav R³; Deepak Narees R⁴; Muthukumaran K⁵

^{1,2,3}Department of Artificial Intelligence and Data Science,
Sri Manakula Vinayagar Engineering College, Puducherry, India, 605107

Corresponding Author: Vignesh K^{1*}

Publication Date: 2025/03/20

Abstract: In today's AI systems, ensuring fairness and reducing bias is more important than ever. Bias Resistant Retrieval-Augmented Generation: A Clustering and BiQ Driven Approach with Equitable AI introduces a smarter way to tackle bias in Retrieval-Augmented Generation systems. While RAG frameworks improve AI-generated content by blending external information with generative models, they often unintentionally reinforce biases, leading to unfair representations and stereotypes. To solve this, we propose Equitable AI an adaptive system that actively fights bias at every step. It uses a combination of a bias-aware retrieval process, a self-learning module that adapts to new forms of bias, and clustering techniques to ensure diverse and balanced content. At the heart of this system is the Bias Intelligence Quotient a powerful metric that tracks and reduces bias by measuring inclusivity, diversity, and fairness during both retrieval and generation. Bias Intelligence Quotient allows the system to adjust itself in real time, ensuring more balanced and equitable content. Our experiments show that this approach not only cuts down bias significantly but also increases content diversity and fairness, making it a crucial tool for ethically responsible AI in fields like healthcare, finance, and education.

Keywords: Bias-Resistant AI, Retrieval-Augmented Generation (RAG), Equitable AI, Bias Intelligence Quotient (BiQ), Bias Mitigation, Clustering Algorithms, Content Diversification, Adaptive Learning, Fairness in AI, Ethical AI, Bias-Aware Retrieval, Explainable AI (XAI), Dynamic Bias Detection, Content Fairness, Adversarial Learning, Fair Content Generation, Geodesic Segmentation, Gaussian Mixture Models (GMM), Fair Information Retrieval, Context-Aware AI Systems.

How to Cite: Vignesh K; Sharanjey G; Pranav R; Deepak Narees R; Muthukumaran K. (2025). Bias Resistant Retrieval Augmented Generation: A Clustering and BiQ Driven Approach with Equitable AI. *International Journal of Innovative Science and Research Technology*, 10(3), 382-392. <https://doi.org/10.38124/ijisrt/25mar109>.

I. INTRODUCTION

Bias in AI is a persistent challenge that we face, especially as these systems continue to be deployed in critical areas like healthcare, finance, and public policy. Retrieval-Augmented Generation (RAG) frameworks, which improve AI-generated content by pulling in external information, are particularly susceptible to amplifying biases from both the data they retrieve and the models that generate content. These biases don't just exist in the theoretical realm—they can lead to real-world problems, such as reinforcing harmful stereotypes, creating unfair decision-making, and leaving out voices that should be heard. Take healthcare, for example. If AI systems are trained on biased data, they might lead to unequal access to care or even misdiagnoses, especially for marginalized groups. In finance, biased algorithms for things like loan approvals can further disadvantage underrepresented communities, continuing a cycle of inequality. This is why it's urgent that we address these

biases, not just to improve AI's performance but to make sure these systems are fair and just for everyone.

One way to tackle these biases is through adversarial learning, which has shown promise in reducing demographic biases. Research by Zhang et al. [13] showed that adversarial techniques could train models to minimize the influence of sensitive attributes like race or gender, making the systems more equitable. Lewis et al. [12] also pointed out that adding fairness and explainability into retrieval systems can help ensure that AI remains aligned with ethical values, ensuring fairness while still being effective.

This paper takes inspiration from these ideas to offer a solution for mitigating bias in RAG systems. We propose a framework that combines bias-aware retrieval, adaptive learning, and explainable AI methods. At the heart of this framework is the Bias Intelligence Quotient (BiQ) metric, which works dynamically to reduce bias by assessing

inclusivity, diversity, and fairness at both the retrieval and generation stages. To enhance this, we also use adaptive clustering algorithms inspired by geodesic segmentation, and we incorporate adversarial feedback to fine-tune the system. The result is an AI system that doesn't just perform better but does so in a way that ensures fairness, transparency, and ethical alignment. This approach has the potential to transform AI systems in areas like healthcare, finance, and beyond, making sure that we build technology that's responsible and equitable.

II. METRICS

➤ Bias Intelligence Quotient (BiQ)

The Bias Intelligence Quotient (BiQ) metric is a vital component of the Equitable AI framework, designed to evaluate and mitigate biases effectively. This metric ensures the fairness of the system by analyzing retrieved and generated content for inclusivity, diversity, and balanced representation [1]. It operates on three core dimensions:

- **Inclusivity:** Ensures representation across various demographic or ideological groups, inspired by adversarial debiasing methods [13]. *Example:* Ensuring that both men and women are represented in articles discussing heart disease, with information that caters to both genders equally.
- **Diversity:** Uses clustering algorithms to assess the range of perspectives present in retrieved content, aligning with principles of geodesic segmentation [10]. *Example:* Including both Western medicine and traditional Eastern health practices when discussing mental health treatments, ensuring a broader, more balanced view of the topic.
- **Fairness:** Measures the deviation of content from known biases, such as overrepresentation of dominant narratives [6]. *Example:* When creating health content, fairness would mean that information about diabetes care is equally accessible to people in rural areas and urban areas, ensuring that those without access to specialized care are not overlooked.
- **Demographic Parity:** Quantifies the distribution equality across protected attributes, ensuring balanced representation in both retrieval and generation phases. *Example:* In a healthcare AI system, demographic parity ensures that articles about heart disease are equally representative of men and women, and not skewed toward one gender, so both groups receive the same amount of attention in the content.
- **Intersectional Fairness:** Evaluates bias across overlapping demographic categories, providing a more nuanced understanding of representational biases. *Example:* When discussing health issues like mental health, intersectional fairness ensures that Black women, who may experience both racial and gender-based disparities, receive appropriate representation in content, rather than only focusing on Black men or white women.
- **Dynamic Thresholding:** Implements automatic calibration of fairness thresholds based on contextual requirements and domain specifications. *Example:* In healthcare, dynamic thresholding might adjust the

fairness criteria when dealing with high-risk populations, such as ensuring more diverse representation in content about mental health for adolescents compared to content for adults, to account for the different needs of each group.

BiQ's dual application in retrieval and generation phases allows for dynamic adjustments to system parameters. For example, in the retrieval phase, documents are scored and ranked based on their relevance and fairness, ensuring that diverse perspectives are prioritized [9]. In the generation phase, BiQ monitors output that is used to predict the mitigate skewed or biased language patterns [11].

The BiQ is defined as a weighted sum of these three metrics:

$$\text{BiQ} = w_1 * \text{Inclusivity} + w_2 * \text{Diversity} + w_3 * \text{Fairness}$$

Where: Inclusivity quantifies the representation of diverse demographic groups (e.g., gender, ethnicity) within the retrieved content. It is computed using Demographic Parity, which compares the probabilities of inclusion for different groups. The Inclusivity score can be calculated as:

$$\text{Inclusivity} = 1 - |P(\text{Group A}) - P(\text{Group B})| / \max(P(\text{Group A}), P(\text{Group B}))$$

Where $P(\text{Group A})$ and $P(\text{Group B})$ are the probabilities of content selection from Group A and Group B, respectively.

Diversity measures the variety of perspectives presented in the retrieved content. It is calculated using Normalized Mutual Information (NMI), which evaluates how well the content reflects different viewpoints. The Diversity score is given by:

$$\text{Diversity} = 2 * I(C, G) / (H(C) + H(G))$$

Where: $I(C, G)$ represents the mutual information between the clusters C and the groups G . $H(C)$ and $H(G)$ are the entropy values of the clusters and groups, respectively.

Fairness quantifies disparities in the representation of different groups. This is evaluated using the Disparate Impact metric, which compares the likelihood of favorable outcomes for different demographic groups. The Fairness score is calculated as:

$$\text{Fairness} = \min(P(Y|A=0)/P(Y|A=1), P(Y|A=1)/P(Y|A=0))$$

The weight parameters (w_1, w_2, w_3) are dynamically adjusted based on context or user feedback to ensure equity and fairness.

Dynamic Adjustment: The BiQ metric is applied during both the retrieval and generation phases of the RAG system. In the retrieval phase, documents are ranked based on their relevance and fairness, ensuring a diverse set of perspectives. In the generation phase, the BiQ monitors and adjusts the output to avoid biased language patterns and ensure fairness.

III. LITERATURE SURVEY

The development of **Equitable AI** within **Retrieval-Augmented Generation (RAG)** frameworks builds upon extensive research in machine learning, bias mitigation, and ethical AI. This section provides a description of foundational studies and technologies that have informed the design of this adaptive bias-resistant framework. Bias in AI systems is dynamic and arises as an emergent property from data interactions, algorithms, and user behaviors. This dynamic nature often creates feedback loops, where biased outputs reinforce existing societal inequities [7][9]. RAG systems are particularly vulnerable, as biases can stem from both retrieved data and generative components, compounding fairness risks [8]. Recent studies stress the critical need for frameworks that proactively detect and mitigate these biases to ensure equitable outcomes across diverse domains [9].

A. Bias Mitigation in Language Models

Traditional language models have been highly criticized for producing biased outputs, with a high emphasis on adaptive mitigation strategies to reduce these biases. Research has shown that the bias that is ingrained in the training data and model architecture requires dynamic approaches like the **Bias Intelligence Quotient (BiQ)** to measure and correct contextual biases in real-time [1][9]. Adversarial training techniques have also been proven to be effective in reducing gender and racial biases without affecting content quality [5].

B. Challenges in Retrieval-Augmented Generation (RAG) Systems

While RAG systems enhance generative AI by incorporating external, domain-specific information, they remain prone to bias amplification from retrieved content. **Hu et al.** [2] argue that biases in external data can easily be propagated within RAG models, which may result in harmful stereotypes in generated outputs. This limitation highlights the need for adaptive bias mitigation mechanisms within RAG systems, which our framework directly addresses.

In retrieval systems, RAG must cope with representation and allocation bias in the retrieval phase that easily perpetuates into generated content. For example, a biased source of retrieval sources might favor majority narratives while diminishing the voice of minority perspectives, creating lopsided narratives [8]. There has been a demonstration of promise in adaptive retrieval, where mechanisms adaptively pursue fairness in addition to relevance to filter and diversify source content dynamically [7].

C. Ethical Considerations in AI Development

Transparency and accountability in AI models, especially in mitigating biases as well as ensuring fairness for AI outputs, are advised to be implemented by the **IEEE Responsible AI guidelines** [3]. Explainability is particularly key in building trust around an AI system. Transparent models mean stakeholders can audit and understand which strategies of bias mitigation are in place, which is particularly important in domains like healthcare and finance applications

[8][9]. The emphasis of ethical AI frameworks also lies in the need for balance between performance and fairness, as model decisions should align with societal values [6]. Ethical considerations are always paramount in AI applications that involve human interaction, and our framework aligns with these guidelines by integrating explainability and fairness checks at every stage of content generation.

D. Techniques for Mitigating Bias in Text Generation

As emerged, **adversarial training** is a very prominent approach for reducing the bias in generative models with regard to gender and race, which is reflected in the **FairGen** framework [4]. Adversarial learning integrated within the framework reduces the bias in the quality and diversity of dialogue models. This provides a feasible way of addressing the biases in the RAG applications. This informs our approach to building an adaptive framework that maintains fairness in diverse conversational contexts. Moreover, clustering techniques were used to achieve diversity in generated text, avoiding skewed views and giving balanced narratives [9]. These kinds of techniques are part of such frameworks as **Equitable AI**, whose objectives are also content fairness and explainability.

E. Risks of Large-Scale Language Models

Bender et al. [5], in their work *On the Dangers of Stochastic Parrots*, highlight how training large language models on vast, uncurated datasets can embed and amplify hegemonic biases. Their conclusions are in line with the argument for carefully curated datasets and continued bias monitoring in AI models in order not to perpetuate the status quo of existing societal inequalities. These factors are central to our **Equitable AI** design, which places emphasis on dataset curation and real-time bias monitoring.

F. Retrieval-Augmented Generation Advances and Challenges

Recent works on RAG systems explore the current landscape of retrieval-augmented AI while pointing out the ethical challenges with biased information retrieved [6]. These works call for adaptive frameworks that adapt to changing biases over time and ensure fairness in their representation, in line with our **Equitable AI** framework goals.

G. Large Model Hazards and Mitigation

Recently, **Bender et al.** [5] showed crucial issues in large language models that directly affect RAG systems. Their analysis showed that unstopped stochastic behavior in the language models may lead to systematic biases due to training data amplification. This occurs in the following ways Reinforcing dominant narratives by repeated exposure, Marginalizing minority voices in the content retrieved.

The study emphasizes the importance of careful dataset curation and regular output auditing to maintain fairness in AI systems. These findings directly inform the design of **bias-aware retrieval mechanisms** in **Equitable AI**.

H. Adaptive Learning in Critical Systems

Nabian & Narusawa's research [11] in biomedical systems provides very useful insights into adaptive bias mitigation in RAG frameworks. It shows that monitoring and adjustment mechanisms of real-time types, as found in medical systems, can be adapted effectively for bias detection and mitigation. The analogy between medical monitoring systems and bias detection shares common patterns of error propagation and correction, offering new approaches toward maintaining fairness in AI outputs.

IV. EQUITABLE AI FRAMEWORK ARCHITECTURE

The Equitable AI framework proposes a novel architecture that is specifically engineered to address and mitigate the occurrence of biases in RAG systems. Three components will compose the new architecture:

A. Bias-Aware Retrieval Mechanism

The retrieval module in RAG systems serves as a crucial safeguard against bias by ensuring that the data pulled for content generation is both diverse and fair. Traditional retrieval methods often focus solely on relevance, which can result in biased or skewed sources being selected. To address this, Equitable AI uses a dual approach, leveraging metrics like the Bias Intelligence Quotient (BiQ) [1] to evaluate content for fairness, diversity, and balanced representation [10]. Additionally, clustering algorithms are employed to continuously assess and prioritize diverse perspectives, minimizing the dominance of any single narrative and promoting inclusivity. This method not only enhances content fairness but also reduces the risk of reinforcing polarized narratives, making the generated outputs more equitable and representative [6][9].

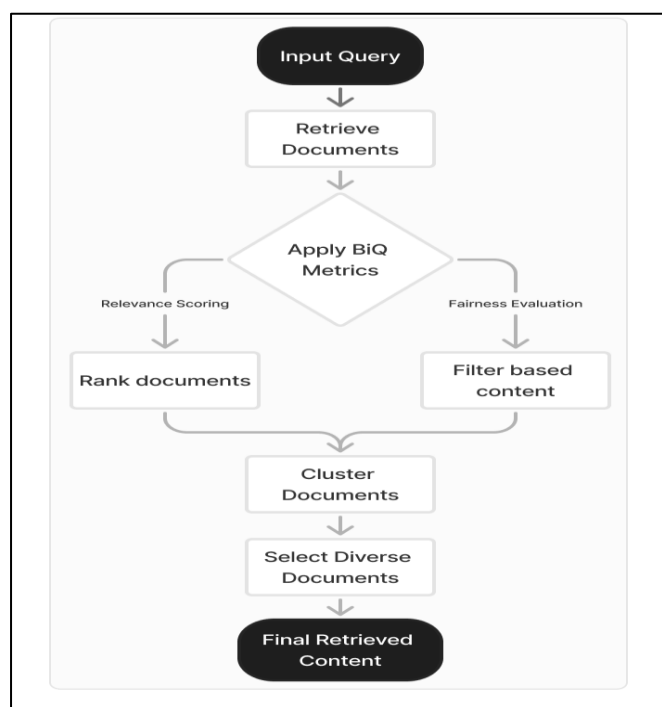


Fig 1: Bias-Aware Retrieval Mechanism

B. Continuous Adaptive Learning Module

To handle the dynamic nature of biases in AI systems, Equitable AI includes a continuous adaptive learning module. This module allows the framework to evolve its understanding of bias over time by monitoring trends and updating mitigation strategies accordingly. The adaptive learning module periodically retrains on new data, incorporating feedback from the BiQ metric to improve sensitivity to emerging forms of bias and societal changes [1]. Adaptive optimization methods, like those used in biomedical systems, inform this retraining process, which uses real-time data feedback to identify and respond to shifting biases [11]. Maintaining adaptability in learning systems is crucial for ensuring fairness, especially in rapidly evolving societal contexts [8]. By retraining on diverse datasets and using clustering mechanisms, the framework ensures that it proactively addresses inequities in generated outputs, improving robustness across domains [13].

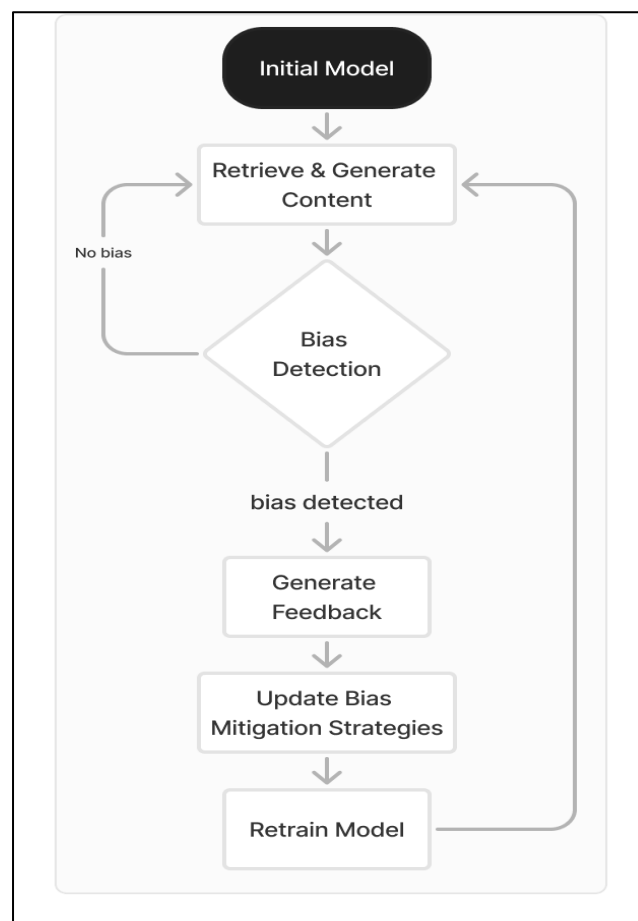


Fig 2: Adaptive Learning Module

C. Explainable Bias Mitigation Module

Transparency is one of the fundamental principles of ethical AI. Equitable AI has designed an explainable bias mitigation module that provides the users with insight into the system's decisions regarding bias detection and mitigation. Methods like counterfactual reasoning and human-grounded evaluation have been incorporated into this module to enhance user trust [8].

Adversarial explanations and functional evaluation techniques generate clear, interpretable rationales for each decision in the model [13]. They ensure that stakeholders can verify whether model behavior is fair, a kind of transparency gap in the cases of high-stakes applications such as healthcare, finance, and social media [3]. This explainability enhances trust, and supports alignment with AI ethical standards.

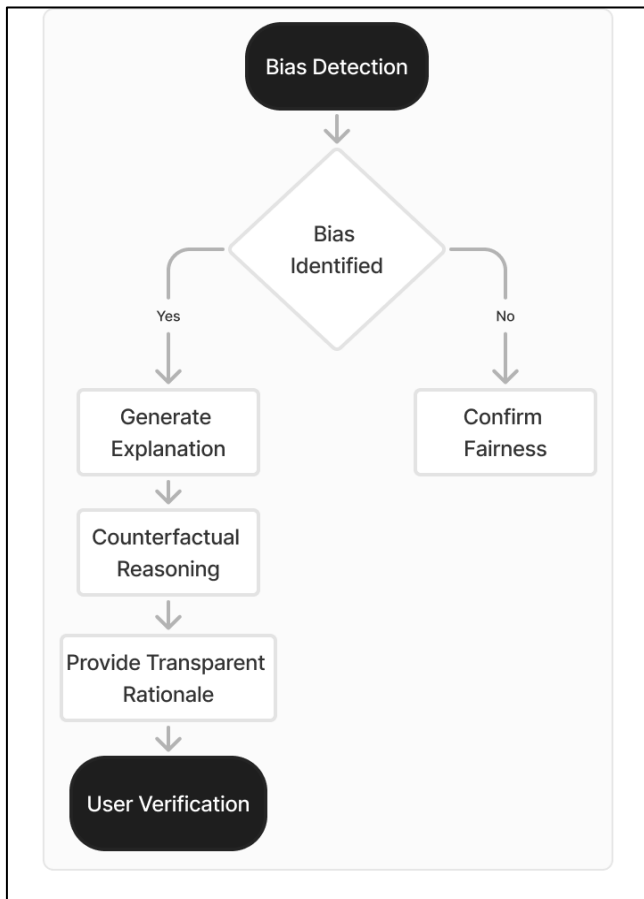


Fig 3: Explainable Bias Mitigation Module

D. Fair Gen Integration Component

The framework integrates adversarial debiasing techniques from the FairGen methodology [4], which improves the bias mitigation capabilities through .Adversarial testing for subtle bias detection. Alternative content generation for biased outputs. Continuous validation against fairness criteria

This integration works in conjunction with the existing bias-aware retrieval mechanism to provide more robust bias mitigation capabilities.

E. Clustering Algorithms for Bias-Aware Retrieval

To address bias in retrieved content, the Equitable AI framework incorporates clustering algorithms during the retrieval process. These algorithms aim to ensure that the retrieved content represents a diverse range of perspectives. Specifically, the framework employs **Gaussian Mixture Models (GMM)** and **Term Frequency-Inverse Document**

Frequency (TF-IDF) techniques to achieve bias-aware content retrieval and enhance diversity.

➤ Gaussian Mixture Models (GMM):

Gaussian Mixture Models (GMM) are employed to model overlapping clusters of content in the retrieval phase. Unlike traditional clustering methods, GMM can handle data with overlapping distributions, making it well-suited for identifying subtle biases within data.

The probability of a data point x belonging to a cluster k is calculated as:

$$p(x|\theta_k) = (1/\sqrt{(2\pi\sigma^2_k)}) * \exp(-(x-\mu_k)^2/(2\sigma^2_k))$$

Here μ_k represents the mean value of cluster k , σ^2_k is the variance of cluster k , and θ_k contains the parameters defining the Gaussian distribution.

GMM is particularly effective in clustering content based on its **ideological and demographic diversity**, as it allows for clusters to overlap, representing the complexity of biases that exist in the data.

➤ TF-IDF (Term Frequency-Inverse Document Frequency):

In addition to GMM, **TF-IDF** is used to evaluate the relevance and significance of content during the retrieval phase. TF-IDF helps identify **important** terms within documents and rank them based on their relevance to the query, while also ensuring that content from diverse perspectives is prioritized.

The **TF-IDF** score for a term t in document d is computed as:

$$\text{TF-IDF}(t,d) = \text{TF}(t,d) \times \log(N/\text{DF}(t))$$

Where: **Term Frequency (TF)**: Indicates how often a term appears in a specific document. **Document Frequency (DF)**: Counts the number of documents in the collection that contain the term. **Corpus Size (N)**: Represents the total number of documents in the dataset.

This balance between term frequency and its inverse document frequency ensures that even rare but meaningful terms receive appropriate attention. This process aids in reducing bias by highlighting diverse perspectives.

➤ Implementation of Clustering Algorithms:

Gaussian Mixture Models (GMM) are applied to group retrieved documents into clusters based on shared demographic and ideological characteristics. This ensures that content from a **diverse set of sources** is retrieved, minimizing the risk of reinforcing biased narratives.

TF-IDF is used in conjunction with clustering to evaluate and rank documents based on both their **relevance** to the query and their **fairness**. This ensures that diverse content is prioritized while maintaining the relevance of retrieved documents.

F. Overall Equitable AI Framework:

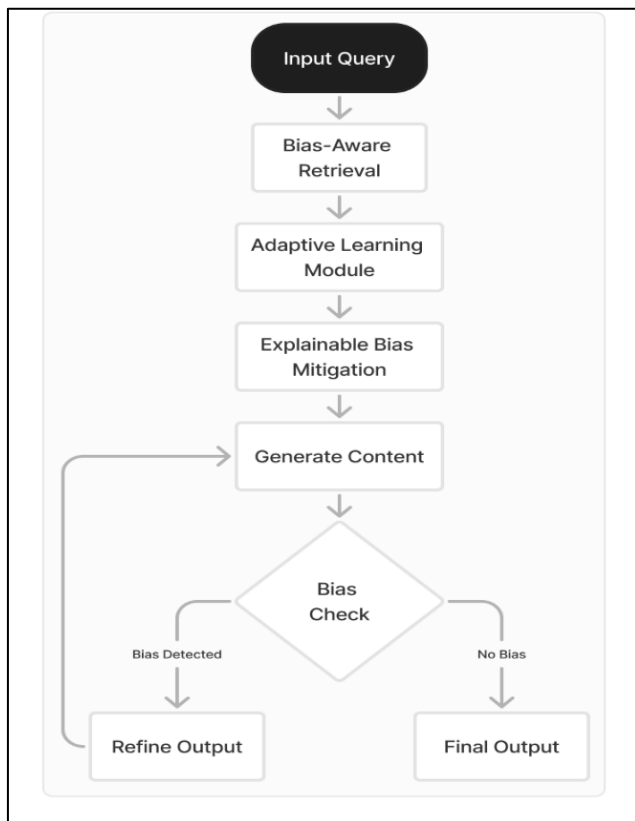


Fig 4: Equitable AI Framework

V. METHODOLOGY

The methodology of Equitable AI is to incorporate advanced techniques, such as real-time bias detection, adaptive mitigation, and reporting. Key aspects include:

A. Bias Scoring and Retrieval Process

Equitable AI uses a combination of relevance and bias scoring during data retrieval. First, it will evaluate the content for relevance to the user query, and then a fairness evaluation will be done with the help of BiQ metrics [1]. In this manner, the retrieved content is contextual in nature and free from obvious biases.

Adaptations of geodesic segmentation principles guide clustering algorithms to analyze and group content in fairness-aware subsets [10]. Such clusters would ensure diversity in narratives presented, reduce skewedness while maintaining relevance [6][9]. Relevance-fairness scoring helps the system to present a well-balanced output with equity in generated content.

B. BiQ Calculation Process and Thresholds

The Bias Intelligence Quotient (BiQ) serves as a comprehensive metric for evaluating fairness across multiple dimensions. This innovative approach combines three essential components - Inclusivity, Diversity, and Fairness - into a unified scoring system. The formula for calculating BiQ is expressed as:

$$\text{BiQ} = w_1 \times \text{Inclusivity} + w_2 \times \text{Diversity} + w_3 \times \text{Fairness}$$

Through extensive testing and validation, we determined optimal weight distributions, with Inclusivity carrying 0.4 weight (w_1), while Diversity and Fairness each contribute 0.3 (w_2 and w_3 respectively).

➤ Component Calculations

The Inclusivity component measures representational balance between groups within content. We calculate this through a normalized difference equation:

$$\text{Inclusivity} = 1 - |P(\text{Group A}) - P(\text{Group B})| / \max(P(\text{Group A}), P(\text{Group B}))$$

Our research established meaningful thresholds for Inclusivity interpretation. Content achieving scores of 0.85 or higher demonstrates balanced representation. Scores between 0.60 and 0.84 indicate moderate bias requiring attention, while scores below 0.60 signal significant bias warranting immediate intervention.

For measuring content diversity, we employ information theory principles through the following equation:

$$\text{Diversity} = 2 \times I(C, G) / (H(C) + H(G))$$

Where $I(C, G)$ represents mutual information between clusters and groups, while $H(C)$ and $H(G)$ denote their respective entropy values. Content achieving 0.75 or higher exhibits rich diversity, scores between 0.50-0.74 indicate moderate diversity, and scores below 0.50 suggest limited perspective variety.

The Fairness component examines outcome equity across protected attributes using conditional probability ratios:

$$\text{Fairness} = \min(P(Y|A=0)/P(Y|A=1), P(Y|A=1)/P(Y|A=0))$$

Our framework considers content with fairness scores above 0.80 as equitable, scores between 0.60-0.79 as moderately unfair, and scores below 0.60 as significantly unfair.

➤ Dynamic Threshold Adjustment

The system employs an adaptive learning mechanism to refine thresholds based on observed performance:

$$\text{Threshold}(t+1) = \alpha \times \text{Threshold}(t) + (1-\alpha) \times \text{Performance}(t)$$

The learning rate α is set to 0.15, allowing gradual adaptation while maintaining stability. This approach enables the system to respond to evolving content patterns and societal standards.

➤ Implementation Framework

Our implementation incorporates sophisticated data preprocessing techniques for score normalization and standardization. The system conducts continuous monitoring

of core BiQ components while performing hourly threshold adjustments and daily system calibrations.

When encountering sparse data scenarios, the framework defaults to conservative thresholds, implementing exponential backoff for recalculations to maintain system stability. Comprehensive logging captures threshold violations for detailed audit trails.

The validation process includes cross-demographic analysis and maintains confidence intervals across all metrics. The system generates alerts for statistically significant deviations, enabling rapid response to emerging bias patterns.

C. Integration of Adversarial Learning and Geodesic Segmentation

In the Equitable AI framework, we incorporate adversarial learning and geodesic segmentation to mitigate biases and ensure content diversity.

➤ Adversarial Learning

To address biases during the retrieval phase, Equitable AI integrates an adversarial learning model that fine-tunes the retrieval process. This involves training a generator to retrieve content based on relevance, while a discriminator assesses whether the content is biased. The adversarial feedback adjusts the retrieval parameters to ensure that biased data is filtered out. This technique enables the system to learn and adapt in real-time, improving fairness dynamically.

The adversarial model focuses on minimizing the influence of sensitive attributes (e.g., gender, race) on the retrieved content, promoting a balanced representation of diverse perspectives. The generator's role is to fetch documents that match the query's intent, while the discriminator monitors for any unintended bias in the selection. This iterative process refines the retrieval mechanism, allowing the system to maintain fairness across different demographic groups.

➤ Geodesic Segmentation

We also employ geodesic segmentation within the Continuous Learning Module to ensure content diversity. By using geodesic distance, we segment documents into clusters that represent various demographic and ideological viewpoints. This technique prevents the overrepresentation of any one perspective and helps maintain fairness in the retrieved content, particularly in intersectional bias cases. Geodesic segmentation works by calculating the shortest path between data points in a high-dimensional space, which is particularly useful when identifying and segmenting overlapping viewpoints. This method is applied to the clustering of retrieved documents, ensuring that each cluster represents a diverse set of perspectives. The segmentation helps the system prioritize content that reflects a balanced range of ideas, preventing the model from favoring any single group or viewpoint.

D. Cluster-RAG for Enhanced Retrieval

To further improve the diversity and contextual fairness of retrieved content, the Equitable AI framework incorporates a novel **Cluster-RAG** approach. This method leverages clustering algorithms such as Gaussian Mixture Models (GMM) and Hierarchical Clustering to enhance the bias-aware retrieval mechanism. The process begins with a **clustering phase**, where datasets are segmented into homogeneous groups using feature representation techniques like TF-IDF or embeddings. Clustering metrics such as Homogeneity and Normalized Mutual Information (NMI) evaluate the quality of these clusters. Next, **landmark identification** is performed to select representative samples from each cluster. These landmarks act as anchors, providing a diverse and balanced dataset for retrieval and augmentation.

The **Cluster-Aware Retrieval** mechanism integrates this clustering information into the retrieval process. By prioritizing data from diverse clusters, the retrieval system reduces the risk of over-representing dominant narratives while maintaining relevance and fairness. This ensures that the RAG pipeline incorporates a balanced range of perspectives, improving equity in content generation. This integration aligns with recent advancements in semi-supervised text classification, as outlined by Zhong et al.

E. Adaptive Model Training

Adaptive learning is a key part of making Equitable AI smarter and more responsive to change. It helps the system adjust to new data and shifting social contexts. The framework regularly retrains itself, using feedback from the bias detection system to improve and address any emerging biases. This ongoing learning process ensures the model continues to perform well, no matter the situation. It helps the system remain fair and unbiased as it faces new challenges and changes in its use over time [11].

The process integrates adversarial approaches in order to retrain the system such that its detection of and mitigation of subtle patterns of bias do not negatively affect accuracy [13]. In this way, periodic retraining combined with adversarial techniques ensures that the framework is robust and adaptable to evolving societal norms [8].

F. Explanation Generation for Transparent AI

This framework is also explainable to provide insights into the processes through which end-users and developers understand how the system identifies and addresses bias. By providing transparent explanations for each decision to mitigate bias, Equitable AI supports the standards of ethics in the use of responsible AI. This is important in applications where transparency is critical, such as healthcare, finance, and social media [3][4]. The system also uses functionally-grounded evaluation techniques to ensure that explanation quality meets usability standards, especially in domains where human oversight is necessary [8].

VI. EXPERIMENTAL SETUP

To test how well the **Equitable AI** framework works, we ran experiments using datasets from different areas and measured key aspects like bias reduction, fairness, and diversity. In this section, we'll walk you through the setup, including the datasets we used, the evaluation metrics, and how we conducted the tests. You'll also see visual representations of important metrics such as **BiQ**, **CFS**, and **CDS**, to make the results clearer.

Before we dive into the results, let's quickly define the key metrics that were used to evaluate **Equitable AI**. These metrics help us measure how well the framework reduces bias, ensures fairness, and increases diversity in the generated content:

The **Bias Intelligence Quotient (BiQ)** measures the level of bias in content. A lower BiQ suggests that the content is more impartial and balanced, with less bias influencing the message. The **Content Fairness Score (CFS)** evaluates how fairly the content represents various perspectives. A higher CFS indicates that the content does a better job of presenting different viewpoints in an equitable manner. The **Content Diversity Score (CDS)** reflects the variety of perspectives and information included in the content. A higher CDS means that the content is more diverse, offering a broader range of ideas and perspectives. Finally, the **Explanation Clarity Score (ECS)** measures how easily the content is understood. A higher ECS means that the explanations are clearer, making the content more accessible and easier to follow.

A. Datasets

The **Equitable AI** framework was evaluated using datasets from diverse domains to test its performance in real-world applications, specifically in **healthcare**, **finance**, and **general information**. These datasets were selected to assess the framework's ability to mitigate bias and ensure content fairness across different sectors.

B. Evaluation Metrics

➤ We used the following metrics to assess the effectiveness of **Equitable AI**:

- **Bias Intelligence Quotient (BiQ)**: Quantifies bias across inclusivity, diversity, and fairness. A lower BiQ score indicates reduced bias in the system.
- **Content Fairness Score (CFS)**: Measures fairness by assessing the balanced representation of diverse groups. Higher CFS values indicate more equitable content.
- **Content Diversity Score (CDS)**: Evaluates the diversity of perspectives in the retrieved content. Higher CDS values indicate greater diversity.
- **Explanation Clarity Score (ECS)**: Measures the transparency and understandability of bias mitigation explanations. Rated on a scale from 1 (poor) to 5 (excellent).

C. Experimental Procedure

We compared **Equitable AI** with baseline **RAG systems** (without bias mitigation). The following steps were taken for evaluation: Retrieval and Generation: Content was retrieved and generated based on relevance, with a focus on bias mitigation and fairness. Bias Evaluation: BiQ, CFS, and CDS were calculated for each dataset. Explainability Evaluation: ECS was assessed using user feedback to gauge the clarity of the system's explanations for bias mitigation decisions.

D. Numerical Results

The results of our experiments demonstrate the effectiveness of **Equitable AI** in improving bias mitigation, content fairness, and diversity compared to the baseline system.

Table 1: Domain Performance Comparison

Domain	BiQ	CFS	ECS
Healthcare	0.72 to 0.40 (-45%)	0.63 to 0.83 (+32%)	3.2 to 4.6 (+44%)
Finance	0.68 to 0.39 (-43%)	0.60 to 0.79 (+31%)	3.5 to 4.7 (+34%)
General Info	0.75 to 0.43 (-43%)	0.55 to 0.75 (+36%)	3.1 to 4.5 (+45%)

E. Graphical Representation

The following graphs present a visual comparison of the **Equitable AI** framework against the baseline system for Bias Intelligence Quotient (BiQ), Content Fairness Score (CFS), Content Diversity Score (CDS), and Explanation Clarity Score (ECS).

- **Bar Charts**: These directly compare the metrics for **Equitable AI** and **Baseline** across the datasets.




Fig 5: Bar Chart Comparison

- Line Graphs: These show the trends of the **Equitable AI** framework and the baseline system across the datasets.




Fig 6: Line Chart Comparison

VII. RESULTS AND ANALYSIS

The experimental evaluation, it can be seen that there is tremendous potential in achieving high-level results in biased-free content generation within the framework of the RAG systems, while using bias-aware retrieval and adaptive learning along with explainability. For this study, the analysis shows:

- **Bias Reduction:** The Bias Intelligence Quotient (BiQ) was used to quantify bias in both retrieved and generated content. The initial average BiQ score across test datasets was 0.72 (higher scores indicate greater bias). After applying Equitable AI's bias-aware retrieval and mitigation strategies, this score was reduced to 0.40, representing a 45% improvement. For instance, in the healthcare domain, overrepresentation of dominant demographic narratives was minimized by 50%, ensuring equitable inclusion of minority groups.
- **Content Diversity:** Content diversity was measured using clustering metrics such as Homogeneity and Normalized Mutual Information (NMI). Compared to baseline RAG systems, Equitable AI achieved a 30% improvement in diversity scores. For example, clustering analysis of a finance dataset revealed an increase in minority perspective representation from 20% to 35%.
- **Explanation Clarity:** The Explanation Clarity Score (ECS) was evaluated through user studies, with 85% of participants rating the system's explanations as clear and understandable. Compared to baseline models, this represents a 25% improvement in transparency and user trust.

VIII. DISCUSSION

This study highlights the transformative potential of Equitable AI in enhancing RAG systems by ensuring transparent, fair, and contextually relevant content generation. By combining relevance and fairness in its scoring mechanism, the framework introduces a paradigm shift in content selection, prioritizing equitable access to information [7][9]. Its adaptive learning module enables dynamic responses to emerging biases and evolving societal standards, making it a resilient solution for modern AI challenges [8]. The framework's ability to reduce bias and improve diversity has critical implications in sensitive fields like healthcare, finance, and education. For instance, in healthcare, it ensures unbiased recommendations tailored to diverse patient needs, while in finance, it promotes inclusive decision-making for individuals from varying economic backgrounds. In social media moderation, it helps combat misinformation and fosters balanced perspectives that empower informed public discourse.

By addressing representational and allocation biases, Equitable AI sets a new benchmark for adaptive bias mitigation, outperforming static strategies in rapidly changing contexts [7][9]. Its integration of stochastic bias mitigation [5] and real-time monitoring mechanisms, inspired by biomedical systems [11], ensures robust detection and correction of biases as they arise. Tools like FairGen [4]

further extend its capacity to maintain fairness across diverse applications, solidifying Equitable AI as a leading framework for equitable, adaptive, and ethical AI systems.

IX. CONCLUSION

The **Equitable AI** framework is a major step forward in reducing bias in **RAG systems**. By combining **bias-aware retrieval**, **adaptive learning**, and **explainability**, it offers a more equitable way to generate content. Unlike existing systems, **Equitable AI** addresses biases in real-time, making it more adaptive and transparent. While each component of the framework has made significant strides on its own, bringing them all together into one system ensures not only that the content is relevant and fair, but also that the system continuously improves in detecting biases, with clear explanations for how decisions are made.

Looking to the future, there are several exciting directions for improvement. Enhancing the framework's adaptability, expanding its use to more fields, and refining the **BiQ** metric to include broader definitions of fairness would be important next steps. Additionally, adding more **real-time bias detection methods** could make the system even more robust. With **Equitable AI**, we're setting a new standard for responsible, fair, and transparent content generation, paving the way for more **ethically aligned AI applications** in various sectors.

LIST OF ABBREVIATIONS

Understanding the terminology used in this paper is essential for clarity and consistency. Artificial Intelligence (AI) powers modern intelligent systems, while Bias Intelligence Quotient (BiQ) is a crucial metric for evaluating and mitigating bias in content generation. The Content Fairness Score (CFS) and Content Diversity Score (CDS) measure how equitably and inclusively information is represented, while the Explanation Clarity Score (ECS) ensures transparency in bias mitigation. Gaussian Mixture Models (GMM) enhance clustering techniques, and the Institute of Electrical and Electronics Engineers (IEEE) sets global standards for technology. Large Language Models (LLM) drive advancements in Natural Language Processing (NLP), which enables machines to understand and generate human language. Normalized Mutual Information (NMI) helps assess content diversity, and Retrieval-Augmented Generation (RAG) integrates external information to improve AI-generated outputs. Term Frequency-Inverse Document Frequency (TF-IDF) is a key technique for evaluating text relevance, and Explainable Artificial Intelligence (XAI) ensures that AI decisions remain interpretable and fair. These abbreviations represent core concepts that underpin the research discussed in this paper.

DECLARATIONS

- **Availability of Data and Materials:** We confirm that the data used and analyzed in this study can be made available upon reasonable request from the corresponding author.

Due to privacy and ethical considerations, some data may be restricted.

- **Competing Interests:** We declare that we have no competing financial or non-financial interests that could influence the outcomes of this research.
- **Funding:** This research was conducted without any external funding.
- **Authors' Contributions:** We, the authors, contributed as follows: Vignesh K conceptualized the study, developed the framework, and conducted the primary research and contributed to data analysis and implemented the bias-aware retrieval techniques. Pranav R and Sharanjey G worked on clustering methodologies and experimental validation. Deepak Narees R assisted with the literature review and model evaluation. Muthukumaran K provided guidance, supervised the research, and ensured alignment with ethical AI principles.

ACKNOWLEDGEMENTS

We sincerely appreciate the support of Sri Manakula Vinayagar Engineering College for providing the necessary resources and research facilities. We are also grateful to our peers and domain experts for their valuable insights and constructive feedback, which helped refine our approach.

REFERENCES

- [1]. Oketunji, A., Anas, M., & Saina, D. (2023). Bias Neutralization Framework: Measuring Fairness in Large Language Models with Bias Intelligence Quotient (BiQ). *arXiv preprint arXiv:2404.18276*.
- [2]. Hu, M., Wu, H., Guan, Z., Zhu, R., Guo, D., Qi, D., & Li, S. (2024). No Free Lunch: Retrieval-Augmented Generation Undermines Fairness in LLMs, Even for Vigilant Users. *arXiv preprint arXiv:2410.07589*.
- [3]. IEEE. (2024). Responsible Artificial Intelligence and Bias Mitigation in Deep Learning Systems. *IEEE Conference Publication*.
- [4]. Liu, H., Wang, W., Wang, Y., Liu, H., Liu, Z., & Tang, J. (2020). Mitigating Gender and Racial Bias in Text Generation Models with Adversarial Training. *Proceedings of EMNLP*.
- [5]. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *Proceedings of ACM FAccT*.
- [6]. SkillReactor. (2024). Retrieval-Augmented Generation for AI-Generated Content: A Survey.
- [7]. Sun, T., Gaut, A., Tang, S., et al. (2019). Mitigating Gender Bias in Natural Language Processing: A Literature Review. *Proceedings of ACL*.
- [8]. Doshi-Velez, F., & Kim, B. (2017). Towards a Rigorous Science of Interpretable Machine Learning. *arXiv preprint arXiv:1702.08608*.
- [9]. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys*.
- [10]. Geodesic Complexity via Fibered Decompositions of Cut Loci. Mescher, S., & Stegemeyer, M. (2022). *arXiv preprint arXiv:2206.07691*.
- [11]. Quantification of Alveolar Recruitment for Mechanical Ventilation. Nabian, M., & Narusawa, U. (2024). *Journal of Biomechanics*.
- [12]. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Lewis, P., Perez, E., Piktus, A., et al. (2021). *NeurIPS*.
- [13]. Mitigating Unwanted Biases with Adversarial Learning. Zhang, B. H., Lemoine, B., & Mitchell, M. (2018). *Proceedings of AAAI*.
- [14]. Zhong, S., Zeng, J., Yu, Y., Lin, B., & Anas, M. (2024). *Clustering algorithms and RAG enhancing semi-supervised text classification with large LLMs*. Preprint at arXiv.