LUISS

Degree Program in Management and Computer Science

# Perception of Artificial Intelligence and How It Differs Based on Different Conditions

**Supervisor**: Prof. Irene Finocchi

**Candidate**: Ayham Alroumi

Academic Year 2024/2025

# Abstract

Artificial Intelligence (AI) is rapidly transforming societies, yet public perceptions of AI remain complex, diverse, and dynamic across different cultures, languages, and demographic groups. This thesis presents a comprehensive, multilingual analysis of online discourse about AI from 2018 to 2025, drawing on a large-scale dataset of articles, social media posts, professional commentaries, and academic content collected from a wide range of sources. Employing advanced sentiment analysis, aspect-based evaluation, and emotion detection techniques, the study systematically examines how attitudes toward AI evolve over time and differ by language, gender, and cultural context. The results reveal that while overall sentiment toward AI hovers near neutrality, significant fluctuations occur in response to major technological and societal events. Notable differences in sentiment and emotional tone are observed across languages and demographic groups, with cultural and thematic factors shaping optimism, skepticism, and key areas of concern such as ethics, employment, and innovation. The research further highlights the methodological challenges of cross-linguistic sentiment analysis and the importance of rigorous data validation. By providing nuanced insights into the drivers of AI perception, this study contributes to a deeper understanding of the societal impact of AI and offers guidance for policymakers, technologists, and communicators seeking to foster responsible and inclusive AI development.

# 1. Introduction

## 1.1 Research Background and Context

Artificial Intelligence (AI) has swiftly evolved from a specific area of computer science to a powerful revolution affecting industries, economies, and daily life. AI technologies are already infusing medicine, financial services, transportation systems, and the arts, among many other fields. This emergence of AI technologies has generated considerable debate, both about the technical features and limitations of AI technologies and about their impacts on society. As AI systems become more obvious and impactful, it is increasingly important from both research and practice perspectives to understand how the AI landscape is viewed by the public, policymakers, and a variety of demographic groups. Especially because public perceptions of AI will be likely to impact adoption, regulation, and patterns of future R&D, attention to public perceptions of AI is an urgent concern in the tech landscape today.

## 1.2 Problem Statement

Despite the prevalence of artificial intelligence (AI), there is still a large gap in our knowledge about how people's perceptions of AI vary by populations, cultures and temporal contexts. Much of the research in this area has been limited to specific areas, languages and demographic groups (often using cross-sectional data that does not truly reflect the changing nature of public attitudes), and has largely missed to capture the differences derived from language, cultural, and media representations informing public perceptions of AI. This limited literature makes it difficult to compare and understand the reception of AI in society, which is a dangerous omission as it also means a lack of understanding of how diverse populations are affected by AI. Such approaches would limit the development of inclusive AI policies and technologies based on the needs, expectations and concerns of diverse communities across the world.

## 1.3 Research Objectives and Questions

The primary objective of this thesis is to provide a comprehensive analysis of public perception of AI across different languages, countries, and demographic groups over the period 2018–2025. Specifically, the research seeks to address the following questions:

1. How does public sentiment toward AI evolve over time across different demographic groups?
2. What are the linguistic and cultural variations in AI perception?
3. How do thematic and emotional dimensions influence public attitudes toward AI?
4. What implications do these perceptions have for AI development, communication, and policy?

## 1.4 Significance of the Study

This study contributes to the field by offering a large-scale, multilingual, and longitudinal analysis of AI perception, addressing critical gaps in previous research. By systematically examining sentiment and emotion across demographic and cultural lines, the research provides insights that are valuable for policymakers, technologists, and scholars seeking to understand and guide the societal impact of AI. The findings have the potential to inform the design of more equitable, culturally sensitive, and publicly accepted AI systems, and to support evidence-based policymaking in an era of rapid technological change.

## 1.5 Scope and Limitations

This study focuses on online public discourse regarding AI from 2018 to 2025, drawing from 4 major languages: English, Spanish, Italian, French; and a range of content platforms, including news media, social media, professional networks and academic sources. While we are presenting a sizable set of data, it is not exhaustive, and certainly does not claim to present all of the global realities and perspectives, particularly from under-represented communities and regions and offline communities. Additionally, we are relying on automated sentiments and machine translation which could have their own biases, as well as being impractical for demographic inference on the data (dependent on the quality and presence of metadata).

Ultimately, while limitations exist, we document these and try to minimize them with the use of methodological rigor and validation when possible.

## 1.6 Thesis Structure Overview

The thesis is organized as follows:

- **Chapter 2** reviews the relevant literature on AI perception, sentiment analysis methodologies, and the theoretical frameworks guiding this research.
- **Chapter 3** details the research design and data collection framework, including the multi-source approach and technical implementation, and presents the methodology for data processing, sentiment analysis, and validation..
- **Chapter 4** reports the results, with sections on temporal, demographic, linguistic, and thematic insights.
- **Chapter 5** concludes with the implications, limitations, and recommendations for future research.

# 2. Literature Review

## 2.1 Theoretical Foundations of AI Perception

Public perceptions of artificial intelligence (AI) are shaped by a complex interplay of technological familiarity, cultural narratives, and personal experiences. As AI becomes embedded in daily life—from healthcare diagnostics to autonomous vehicles—its societal impact is increasingly framed through dual lenses of optimism and apprehension. Theoretical frameworks like the Technology Acceptance Model (TAM) and Value-Based Acceptance Model

(VAM) suggest that perceived usefulness and ethical alignment drive adoption, yet these models often fail to account for the emotional and cultural dimensions of AI acceptance.

A critical theoretical divide emerges between instrumental views (AI as a tool for efficiency) and relational views (AI as a social actor), with collectivist cultures more likely to embrace the latter. For instance, Chinese respondents in cross-cultural studies perceive AI as an extension of communal identity, whereas U.S. participants emphasize individual autonomy and control. This dichotomy underscores the need for culturally grounded theories of AI perception, particularly as global AI development remains dominated by Western epistemological frameworks.

## 2.2 Empirical Insights into Public Attitudes

Recent empirical studies reveal three consistent patterns in AI perception:

1. **Ambivalence Toward Benefits vs. Risks**
   While 52% of Americans express concern about AI's societal impact, 57% welcome its use for routine tasks like household chores. Similarly, 51% of those concerned about racial bias in healthcare believe AI could mitigate discrimination, yet 83% fear job losses from autonomous vehicles. This paradox reflects what Schwesig et al. term the "dual-edge heuristic," where the public simultaneously acknowledges AI's potential and its disruptive consequences.
2. **Demographic Variations**
   *Research consistently demonstrates that demographic factors such as gender and age play a significant role in shaping attitudes toward AI. Surveys and empirical studies indicate that men are generally more enthusiastic and optimistic about AI, while women express higher levels of skepticism and anxiety. For example, one recent survey found that 51% of men described themselves as enthusiastic about AI, compared to 40% of women. Age is also influential, with younger adults tending to be more positive and trusting of AI than older generations, as seen in the contrast between Gen Z and Baby Boomer respondents. However, some recent research suggests that the effect of age may be diminishing as AI becomes more integrated into everyday life. Education and*

*technological literacy are also recognized as important, though findings are more variable and context-dependent*.

3. **Cultural and Geopolitical Divides**

   Adoption rates and public sentiment toward AI vary markedly across regions. For example, recent international surveys show that 83% of Chinese respondents, 80% of Indonesians, and 77% of Thais believe AI products and services offer more benefits than drawbacks, whereas only 40% of Canadians, 39% of Americans, and 36% of Dutch respondents share this view. Similarly, at the organizational level, AI adoption is reported by 50–59% of companies in China, India, and Singapore, compared to just 26–33% in France, Spain, and the United States. These disparities reflect deeper cultural patterns: collectivist societies such as those in East and Southeast Asia tend to view AI as a communal benefit, while individualistic societies in Western Europe and North America are more likely to emphasize concerns about autonomy and risk.

## 2.3 Factors Shaping AI Perception

### 2.3.1 Demographic Influences

Women face disproportionate AI job risks (79% vs. 58% of men) and express greater skepticism (53% restrict children's AI use vs. 26% of men). 74% of adults over 50 distrust AI in healthcare, compared to 12% of 18–29-year-olds. 74% of college graduates support AI regulation, vs. 70% of non-degree holders. Only 30% of U.S. adults can identify common AI applications.

### 2.3.2 Cultural and Linguistic Dimensions

Individualistic Cultures: 39% of Americans view AI as beneficial, vs. 83% of Chinese ([Stanford HAI AI Index, 2025]). Language: Arabic/Mandarin speakers associate AI with collaboration; English speakers emphasize automation ([arXiv, 2024]).

### 2.3.3 Temporal Dynamics .

AI ethics scandals reduced public trust by 22% within months ([Stanford HAI AI Index, 2025]).

## 2.4 Methodological Advances in Sentiment Analysis

Recent advances in sentiment analysis, driven by natural language processing and machine learning, have enabled the systematic classification of attitudes and emotions in large-scale text data. The most common approach, polarity scoring, categorizes text as positive, negative, or neutral, but often struggles with cultural nuances, sarcasm, and domain-specific language. Aspect-based sentiment analysis (ABSA) offers more targeted insights by linking sentiment to specific topics or features, allowing researchers to identify which aspects of AI—such as ethics or healthcare—are viewed positively or negatively. Emotion detection further enriches sentiment analysis by identifying distinct feelings like hope, fear, or joy within text, providing a deeper understanding of public discourse.

Despite these methodological advances, significant challenges persist. Sentiment analysis across languages and cultures is complicated by idiomatic expressions, regional dialects, and translation bias, which can skew results and reduce accuracy. Even advanced models may misinterpret mixed or ambiguous sentiment, highlighting the need for careful validation and culturally sensitive approaches in sentiment research.

## 2.5 Critical Research Gaps

Despite the rapid growth of research on public perceptions of artificial intelligence, several important gaps remain.

First, there is a clear geographical imbalance: the majority of studies focus on North America and Europe, with relatively little attention paid to perspectives from the Global South or other underrepresented regions (Stanford HAI AI Index 2025; Carnegie Endowment 2020). This limits the generalizability of findings and may overlook culturally specific attitudes toward AI.

Second, longitudinal research on AI perception is rare. Most studies rely on cross-sectional surveys or single-point-in-time analyses, making it difficult to track how attitudes evolve as AI technologies become more prevalent or as major events unfold (Frontiers in Computer Science, 2023; Esade/Current Opinion in Psychology, 2025).

Third, there is a persistent gap between public expectations and current AI capabilities, particularly in areas such as empathy and ethical reasoning. While the public often anticipates human-like understanding or emotional intelligence from AI, existing systems frequently fall short, contributing to a "trust gap" and potential disillusionment (Esade/Current Opinion in Psychology, 2025; Stanford HAI AI Index 2025).

Finally, media literacy deficits present a significant challenge: recent surveys indicate that only about one-third of Americans feel confident distinguishing AI-generated content from human-created content, suggesting that many people may misjudge the authenticity or risks associated with AI-driven media (Pew Research, 2023; Stanford HAI AI Index 2025).

## 2.6 Theoretical Framework for the Current Study

This research adopts a **Multidimensional Perception Model (MPM)** that integrates four core theoretical lenses:

1. Temporal Dynamics
   Tracks sentiment evolution through event-based analysis, capturing shifts in public opinion linked to technological breakthroughs (e.g., ChatGPT release).
2. Demographic Mediation
   Accounts for gender and language acknowledging their documented influence on trust, risk perception, and adoption rates.
3. Sentiment-Cognition Interplay
   Combines polarity scoring with aspect-based analysis to capture nuanced attitudes, enabling simultaneous assessment of overall sentiment and topic-specific evaluations.

The MPM addresses key research gaps by:

- Prioritizing cross-linguistic comparisons across six target languages (English, Spanish, Italian, French) to uncover cultural and regional variations.
- Employing longitudinal data (2018–2025) to map perception trajectories and identify persistent vs. transient trends.
- Validating machine-translated content through native speaker audits to mitigate translation bias and ensure linguistic fidelity.

# 3. Methodology

## 3.1 Research Design and Approach

### Overview of the Data Collection System

The data collection system developed for this research represents a sophisticated, multi-layered architecture designed to gather, process, and analyze web content related to artificial intelligence perceptions across different temporal, demographic, and linguistic dimensions. The system implements a robust pipeline that combines automated web scraping, content extraction, and intelligent processing capabilities to create a comprehensive dataset for sentiment analysis.

At its core, the system employs a modular architecture that separates concerns into distinct components, each responsible for specific aspects of the data collection process. The primary data collection engine utilizes the Google Custom Search API as its foundation, enabling systematic and controlled access to web content while adhering to API usage policies and rate limits. This approach ensures reliable and consistent data collection while maintaining compliance with service provider terms of use.

The system implements sophisticated content extraction mechanisms that can handle various types of web content, from simple static pages to complex JavaScript-heavy applications. This is achieved through a dual-mode content extraction strategy: a JavaScript-enabled mode using Playwright with stealth capabilities for dynamic content, and a lightweight mode for static content that doesn't require JavaScript execution. The system intelligently determines which mode to use based on an analysis of the target webpage's characteristics, optimizing resource usage while ensuring comprehensive content capture.

Content processing is handled through a series of specialized parsers, each designed to handle specific types of web content. These include dedicated parsers for major platforms such as Reddit, LinkedIn, Quora, YouTube, and PDF documents. Each parser implements platform-specific logic to extract relevant content while maintaining the context and structure of the original material. The system employs advanced text extraction techniques, including the use of the MainContentExtractor class, which intelligently identifies and extracts the main content from web pages while filtering out navigation elements, advertisements, and other non-essential content.

The system implements robust error handling and recovery mechanisms throughout the data collection pipeline. This includes automatic retry logic for failed requests, rate limiting to prevent API throttling, and comprehensive logging of all operations. The logging system captures detailed information about each step of the process, enabling debugging and monitoring of the data collection process. This is particularly important for long-running collection tasks that may span multiple days or weeks.

Data storage and management are handled through a flexible system that supports multiple output formats, including CSV and JSON. The system implements checkpointing mechanisms that allow for the resumption of interrupted collection tasks, ensuring that no data is lost in case of system failures or interruptions. This is particularly important for large-scale data collection operations that may need to be paused and resumed multiple times.

The system also implements sophisticated content deduplication mechanisms to ensure the quality and uniqueness of the collected data. This includes both exact matching and fuzzy matching algorithms to identify and remove duplicate content, even when it appears in slightly different forms across different sources. The deduplication process takes into account various factors such as content similarity, publication dates, and source credibility.

To ensure the quality and relevance of the collected data, the system implements various validation and filtering mechanisms. These include checks for content relevance, language detection, and basic quality metrics. The system can also handle multilingual content through integration with translation services, enabling the collection and analysis of content in multiple languages.

The data collection system is designed to be scalable and maintainable, with clear separation of concerns and modular architecture. This allows for easy extension and modification of the system's capabilities as new requirements emerge or as the research scope evolves. The system's architecture also enables parallel processing of data collection tasks, significantly improving the efficiency of the collection process.

## Multi-source Approach

The research implements a comprehensive multi-source data collection strategy that leverages multiple platforms and content types to gather diverse perspectives on artificial intelligence. This approach ensures a rich and varied dataset that captures different aspects of public perception and discourse surrounding AI. The system's multi-source architecture is designed to handle various content formats and platforms while maintaining consistent data quality and structure.

### Social Media and Discussion Platforms

The system incorporates specialized parsers for major social media and discussion platforms, each implementing platform-specific extraction logic to capture the unique characteristics of content from these sources. Reddit content collection, for instance, utilizes the official Reddit API through the PRAW (Python Reddit API Wrapper) library, enabling access to both post content and associated comments. The Reddit parser implements sophisticated comment threading analysis, allowing for the collection of not just the main post content but also the most relevant and engaging discussions in the comments section. This is particularly valuable for capturing nuanced public opinions and debates about AI topics.

Quora content extraction employs a specialized approach that leverages the platform's structured data format. The parser extracts content from Quora's application/ld+json scripts, which contain rich metadata about questions, answers, and their authors. This approach allows for the collection of detailed information about the context of discussions, including the expertise level of respondents and the temporal aspects of the content. The Quora parser implements sophisticated

JSON extraction and parsing logic to handle the platform's complex data structure while maintaining the hierarchical relationship between questions and answers.

### Professional and Academic Content

The system includes dedicated parsers for professional and academic content sources. LinkedIn content collection focuses on professional perspectives and industry discussions about AI, capturing insights from technology professionals, researchers, and industry leaders. The parser is designed to handle LinkedIn's dynamic content loading and authentication requirements, ensuring reliable access to professional discussions and articles.

For academic and technical content, the system implements a PDF parser that can extract and process content from academic papers, technical documentation, and research reports. This parser employs advanced text extraction techniques to handle various PDF formats while preserving the structure and formatting of the original documents. The PDF parser is particularly valuable for capturing formal academic discourse and technical perspectives on AI development and implementation.

### Video Content and Multimedia

YouTube content collection is handled through a specialized parser that extracts both video metadata and associated comments. This parser is designed to capture not just the video content itself but also the surrounding discussion and community engagement. The YouTube parser implements sophisticated comment threading and moderation analysis, allowing for the collection of relevant discussions and reactions to AI-related video content.

**Content Normalization and Integration**

The multi-source approach includes a robust content integration layer that normalizes data from different sources into a consistent format. This layer implements various processing steps to ensure that content from different platforms can be effectively compared and analyzed. The integration process includes:

- Content structure normalization to ensure consistent formatting across different sources
- Metadata standardization to maintain uniform information about content origin, date, and context
- Text cleaning and normalization to handle different writing styles and formats
- Language detection and processing to support multilingual content analysis
- Content deduplication across different sources to maintain data quality

**Platform-Specific Considerations**

Each platform parser implements specific handling for platform-unique features and limitations:

1. Rate limiting and API quota management to comply with platform restrictions
2. Authentication and access token management for platforms requiring user authentication
3. Content structure preservation to maintain the context and meaning of the original content
4. Error handling and recovery mechanisms specific to each platform's API behavior
5. Content validation to ensure compliance with platform terms of service

This sophisticated multi-source architecture enables comprehensive capture of AI perception signals across professional, academic, and public discourse channels. By maintaining platform-specific parsing excellence while enforcing unified quality standards, the system ensures the collected dataset's validity for comparative analysis across demographic, temporal, and linguistic dimensions.

## 3.2 Data Collection Framework

### 3.2.1 Core Component

#### Web Search Engine Integration

The web search engine integration component represents a sophisticated implementation that leverages the Google Custom Search API to systematically collect and process web content related to artificial intelligence perceptions. This component is built around the AIPerceptionTimeline class, which implements a robust and efficient system for interacting with Google's search infrastructure while adhering to API limitations and best practices.

The integration begins with a carefully designed initialization process that establishes a secure connection to the Google Custom Search API. The implementation uses the official Google API Python client library, which provides a reliable and well-maintained interface to the search service. The initialization process requires two critical components: a Google API key and a Custom Search Engine ID (cx). These credentials are managed securely and are essential for authenticating requests to the Google Custom Search API. The API key is used to authenticate the application, while the Custom Search Engine ID defines the scope and configuration of the search, allowing for customized search behavior and result filtering.

The core of the search integration lies in the search_for_year method, which implements a sophisticated approach to retrieving search results for specific time periods. This method is designed to handle the limitations of the Google Custom Search API, which restricts the number of results that can be retrieved in a single request. The implementation employs a pagination strategy that makes multiple API calls when necessary to collect the desired number of results. The method constructs year-specific search queries using Google's date range operators, allowing for precise temporal filtering of results. This is achieved by appending "after:" and "before:" parameters to the search query, creating a one-year window around the target year.

The search result processing includes comprehensive error handling and rate limiting mechanisms. The implementation includes sophisticated error detection for various HTTP status codes, with special handling for rate limit errors (status codes 403 and 429). When rate limits are encountered, the system implements an exponential backoff strategy, waiting for an appropriate duration before retrying the request. This approach ensures reliable operation while respecting Google's API usage policies and preventing service disruption.

The search integration includes a robust result processing pipeline that extracts and normalizes various pieces of information from the search results. For each result, the system captures the title, domain, URL, and snippet (preview text) from the search result. The implementation includes sophisticated URL parsing to extract the domain from the full URL, which is crucial for categorizing and analyzing the sources of the content. The result processing also includes data validation and cleaning steps to ensure the quality and consistency of the collected data.

The integration implements a sophisticated progress tracking and reporting system that provides real-time feedback on the collection process. This is achieved through the use of the Rich library, which provides a rich terminal interface with progress bars, status updates, and detailed logging. The progress tracking system includes features such as elapsed time, remaining time, and percentage completion, giving users a clear understanding of the collection process's status and progress.

The search integration includes comprehensive data persistence mechanisms that allow for saving the collected results in multiple formats. The implementation supports both CSV and JSON output formats, with automatic filename generation based on timestamps to prevent overwriting of existing data. The data persistence system includes validation checks to ensure that data is only saved when results have been successfully collected, preventing the creation of empty or incomplete data files.

The integration also includes a sophisticated resumption capability that allows for continuing interrupted collection processes. This is implemented through the resume_from parameter in the collect_timeline method, which enables the system to skip already processed years and continue from a specific point in the timeline. This feature is particularly valuable for long-running collection tasks that may need to be interrupted and resumed due to various factors such as network issues or API quota limitations.

The web search engine integration component is designed to be both robust and flexible, allowing for customization of various parameters such as the number of results per year, delay between requests, and the time range for collection. The implementation includes comprehensive input validation to ensure that all parameters are within acceptable ranges and that the system operates within the constraints of the Google Custom Search API. This validation includes checks for empty queries, valid year ranges, and appropriate delay values, ensuring that the system operates reliably and efficiently.

### Temporal Data Collection

The temporal data collection component represents a sophisticated implementation designed to systematically gather and organize web content across a defined time period, specifically focusing on the evolution of AI perception from 2018 to 2025. This component is implemented through the collect_timeline method, which orchestrates a complex process of year-by-year data collection while maintaining data integrity and temporal consistency.

The temporal collection process begins with a comprehensive input validation system that ensures the integrity of the collection parameters. The implementation includes rigorous validation of the start year, end year, and results per year parameters, enforcing strict boundaries to ensure the collection process remains within the scope of the research. The validation system checks for logical consistency in the year range, ensuring that the start year is not after the end year and that both years fall within the acceptable range of 2018 to 2025. This validation is crucial for maintaining the temporal integrity of the collected data and preventing invalid collection attempts.

The core of the temporal collection process lies in its year-by-year processing approach. The implementation creates a chronological sequence of years to process, maintaining the temporal order of data collection. This chronological processing is essential for understanding the evolution of AI perception over time and ensures that the data can be properly analyzed in a temporal context. The system processes each year sequentially, collecting the specified number of results for each year while maintaining proper spacing between collection attempts to respect API rate limits.

The temporal collection includes a sophisticated resumption capability that allows for continuing interrupted collection processes. This feature is particularly valuable for long-running collection tasks that may need to be paused and resumed due to various factors such as network issues, API quota limitations, or system maintenance. The resumption system maintains the temporal integrity of the collection process by allowing it to continue from any point in the timeline without duplicating or missing data. This is achieved through the resume_from parameter, which enables the system to skip already processed years and continue from a specific point in the timeline.

The implementation includes a comprehensive progress tracking system that provides real-time feedback on the collection process. This system uses the Rich library to create a rich terminal interface that displays detailed information about the collection progress, including the current year being processed, the percentage of completion, elapsed time, and estimated remaining time. This progress tracking is crucial for monitoring long-running collection tasks and ensuring that the process is proceeding as expected.

The temporal collection process implements sophisticated error handling and recovery mechanisms. The system includes comprehensive error detection for various failure scenarios, with special handling for API rate limits and network issues. When errors occur, the system implements appropriate recovery strategies, such as exponential backoff for rate limit errors, ensuring that the collection process can continue reliably even in the face of temporary failures.

The temporal collection includes a sophisticated data organization system that maintains the temporal context of the collected data. Each collected result is tagged with its corresponding year, allowing for proper temporal analysis of the data. The system also maintains metadata about the collection process, including timestamps and collection parameters, which is crucial for understanding the context and limitations of the collected data.

The implementation includes comprehensive data persistence mechanisms that allow for saving the collected results in multiple formats. The system supports both CSV and JSON output formats, with automatic filename generation based on timestamps to prevent overwriting of existing data. The data persistence system includes validation checks to ensure that data is only saved when results have been successfully collected, preventing the creation of empty or incomplete data files.

The temporal collection process includes sophisticated rate limiting and delay mechanisms to ensure reliable operation while respecting API usage policies. The system implements configurable delays between collection attempts for different years, allowing for fine-tuning of the collection process based on API limitations and requirements. This rate limiting is crucial for maintaining reliable access to the search API and preventing service disruption.

The temporal collection component is designed to be both robust and flexible, allowing for customization of various parameters such as the number of results per year, delay between requests, and the time range for collection. The implementation includes comprehensive input validation to ensure that all parameters are within acceptable ranges and that the system operates within the constraints of the search API. This validation includes checks for empty queries, valid year ranges, and appropriate delay values, ensuring that the system operates reliably and efficiently.

## The Multi-source Approach

The research implements a sophisticated multi-platform data collection framework that systematically aggregates and processes diverse content types through specialized parsing implementations. This architecture combines web search capabilities with platform-specific content extraction mechanisms, creating a robust pipeline for capturing heterogeneous perspectives on artificial intelligence across temporal, demographic, and linguistic dimensions.

### Social Media and Discussion Platform Parsers

The system incorporates specialized parsers for major social media and discussion platforms, each implementing platform-specific extraction logic to capture the unique characteristics of content from these sources.

#### Reddit Content Integration

The Reddit content parser represents a sophisticated implementation that leverages the official Reddit API through the PRAW (Python Reddit API Wrapper) library. This parser is designed to handle the complex hierarchical structure of Reddit content, which includes posts, comments, and nested discussions. The implementation begins with a robust URL parsing system that can handle various Reddit URL formats, including standard post URLs, shortened URLs, and direct comment links. The parser employs regular expression patterns to extract post IDs from these different URL formats, ensuring reliable content access regardless of how the URL is structured.

Once the post ID is extracted, the parser establishes a connection to the Reddit API using secure authentication credentials. The implementation includes sophisticated comment threading analysis that can handle Reddit's "More Comments" feature, which is used to load additional comments in long threads. The parser is configured to replace these placeholders with actual comments, up to a depth of three levels, ensuring comprehensive coverage of the discussion while maintaining reasonable processing times.

The content extraction process is designed to preserve the hierarchical structure of Reddit discussions. The parser captures the post title, main content, and up to ten top-level comments, maintaining the context and flow of the conversation. This approach allows for the collection of not just the main post content but also the most relevant and engaging discussions in the comments section, which often contain valuable insights and diverse perspectives on AI-related topics.

**Quora Knowledge Extraction**

The Quora content parser implements a unique approach that leverages the platform's structured data format through its application/ld+json scripts. This implementation is particularly sophisticated as it works with Quora's semantic markup, which provides rich metadata about questions, answers, and their authors. The parser begins with a robust HTTP request system that handles various response encodings and potential network issues.

The core of the Quora parser lies in its JSON extraction and parsing logic. The implementation uses regular expressions to locate and extract the application/ld+json script from the HTML content, which contains structured data about the question and its answers. The parser then employs sophisticated JSON parsing to handle the complex nested structure of Quora's data format, which includes information about the main question, suggested answers, and accepted answers.

The content formatting process is particularly detailed, handling various edge cases in Quora's data structure. The parser can process both single answers and lists of answers, and it maintains the hierarchical relationship between questions and their responses. The implementation includes comprehensive error handling for various scenarios, such as missing data fields or malformed JSON content. The output formatting preserves important metadata such as author names, dates, and the structure of the answers, ensuring that the context and credibility of the content are maintained.

### Professional and Academic Content Processors

#### LinkedIn Professional Discourse Analysis

The LinkedIn content parser is designed to handle the professional networking platform's unique content structure and authentication requirements. The implementation focuses on extracting content from LinkedIn posts while maintaining the professional context of the discussions. The parser employs a sophisticated meta-description extraction approach that can handle various HTML structures and content formats.

The parser's implementation includes multiple layers of content extraction to ensure reliable data collection. It first attempts to extract content using a primary meta-description pattern, and if that fails, it falls back to an alternative pattern. This redundancy ensures robust content extraction even when LinkedIn's HTML structure varies. The parser includes HTML entity unescaping to properly handle special characters and formatting in the extracted content.

The implementation also includes comprehensive error handling for various scenarios, such as network issues, authentication failures, or changes in LinkedIn's HTML structure. The parser maintains the professional context of the content by preserving the structure and formatting of the original posts, which is crucial for maintaining the credibility and context of professional discussions about AI.

**PDF Document Processing**

The PDF document parser is designed to handle the complex structure of PDF files, which can contain various types of content and formatting. The implementation uses the PyPDF2 library to extract text content while preserving the document's structure and formatting. The parser begins with a robust file download system that can handle various PDF formats and sizes.

The core of the PDF parser lies in its text extraction capabilities, which are implemented using PyPDF2's sophisticated page parsing features. The parser processes each page of the PDF document individually, extracting text while maintaining the document's structure. The implementation includes comprehensive error handling for various scenarios, such as corrupted PDF files, password-protected documents, or network issues.

The parser's output formatting is designed to maintain the document's structure while making it suitable for text analysis. The implementation preserves paragraph breaks and formatting, which is crucial for maintaining the context and meaning of the content. The parser also includes memory-efficient processing of large PDF files through streaming download and processing.

### Multimedia Content Integration

**YouTube Video Content Analysis**

The YouTube content parser is designed to handle the complex nature of video content and its associated metadata. The implementation focuses on extracting both video transcripts and associated metadata, providing a comprehensive view of the video content. The parser begins with a sophisticated URL parsing system that can handle various YouTube URL formats, including standard watch URLs, shortened URLs, embed URLs, and legacy URL formats.

The core of the YouTube parser lies in its transcript extraction capabilities, which are implemented using the YouTube Transcript API. The parser can handle various transcript formats and languages, and it includes sophisticated error handling for cases where transcripts are disabled or unavailable. The implementation preserves the temporal structure of the video content by maintaining the timing information in the transcript data.

The parser's output formatting is designed to maintain the flow and context of the video content while making it suitable for text analysis. The implementation includes comprehensive error handling for various scenarios, such as unavailable transcripts, network issues, or changes in YouTube's API structure. The parser also maintains metadata about the video, such as its title and description, which provides important context for the content analysis.

Each of these parsers is integrated into a unified data collection system that maintains consistent data quality and structure across different content types. The system includes sophisticated error handling, rate limiting, and content validation mechanisms to ensure reliable data collection while respecting platform terms of service and API limitations. The implementation of these parsers demonstrates a deep understanding of each platform's unique characteristics and challenges, enabling comprehensive and reliable data collection for AI perception analysis.

### 3.2.2 Data Collection Process

#### Search Query Construction

The search query construction process represents a sophisticated implementation designed to generate precise and effective search queries that capture the evolution of AI perception over time. The implementation begins with a base query that focuses on AI perception and public opinion, which is then enhanced with temporal parameters to create year-specific search queries. The query construction process is implemented through the search_for_year method, which combines the base query with Google's date range operators to create precise temporal filters.

The query construction process employs a sophisticated approach to temporal filtering, using Google's "after:" and "before:" operators to create a one-year window around the target year. This approach ensures that the search results are properly contextualized within their temporal period while allowing for some flexibility in the exact publication dates. The implementation constructs the year-specific query by appending these temporal operators to the base query, creating a search string that effectively captures content from the target year.

The query construction process includes comprehensive validation of the search parameters to ensure the effectiveness of the search queries. The implementation checks for empty queries and validates the temporal parameters to ensure they fall within the acceptable range of 1995 to 2025. This validation is crucial for maintaining the quality and relevance of the search results and preventing invalid search attempts.

The query construction process is designed to be flexible and adaptable, allowing for customization of the search parameters based on specific research needs. The implementation supports various query formats and can be easily modified to include additional search parameters or filters. This flexibility is particularly valuable for adapting the search strategy based on the characteristics of different time periods or specific aspects of AI perception that are being investigated.

### Content Extraction

The content extraction process represents a sophisticated implementation designed to extract and process content from various web sources while maintaining the context and structure of the original material. The implementation begins with a robust URL parsing system that can handle various URL formats and extract the necessary identifiers for content retrieval. The content extraction process is implemented through a series of specialized parsers, each designed to handle specific types of web content.

The content extraction process includes sophisticated error handling and recovery mechanisms to ensure reliable operation even in the face of various failure scenarios. The implementation includes comprehensive error detection for various issues such as network problems, API rate limits, or content accessibility issues. When errors occur, the system implements appropriate recovery strategies, such as exponential backoff for rate limit errors or alternative content extraction methods for inaccessible content.

The content extraction process implements a sophisticated approach to handling different content types and formats. The implementation includes specialized parsers for various content types, such as HTML pages, PDF documents, or video transcripts, each implementing specific logic to extract and process the content while maintaining its structure and context. This approach ensures that the extracted content is properly formatted and suitable for subsequent analysis.

The content extraction process includes comprehensive content validation and cleaning mechanisms to ensure the quality and consistency of the extracted content. The implementation includes various validation steps that check for common issues such as missing content, malformed data, or incomplete information. The process also includes content cleaning steps that remove irrelevant elements such as advertisements, navigation menus, or other non-essential content while preserving the important context and structure of the original material.

### Data Storage and Management

The data storage and management process represents a sophisticated implementation designed to store and organize the collected data in a structured and accessible format. The implementation includes comprehensive data persistence mechanisms that support multiple output formats, including CSV and JSON, with automatic filename generation based on timestamps to prevent overwriting of existing data.

The data storage process includes sophisticated data organization mechanisms that maintain the temporal and contextual structure of the collected data. The implementation preserves important metadata such as publication dates, source information, and content types, which is crucial for maintaining the context and credibility of the collected data. The process also includes data validation checks to ensure that data is only saved when results have been successfully collected, preventing the creation of empty or incomplete data files.

The data management process includes comprehensive error handling and recovery mechanisms to ensure reliable operation even in the face of various failure scenarios. The implementation includes sophisticated error detection for various issues such as file system problems, permission issues, or data corruption. When errors occur, the system implements appropriate recovery strategies, such as automatic retry mechanisms or alternative storage methods.

The data management process implements a sophisticated approach to data organization and retrieval. The implementation includes mechanisms for organizing the collected data by various criteria such as temporal period, content type, or source, making it easy to access and analyze specific subsets of the data. The process also includes comprehensive logging and monitoring mechanisms that track the status and progress of the data collection and storage operations, providing valuable insights into the system's operation and performance.

## 3.3 Technical Implementation

### 3.3.1 System Architecture

The system architecture represents a sophisticated, modular design that implements a comprehensive data collection and processing pipeline for analyzing AI perception across different temporal, demographic, and linguistic dimensions. The architecture is organized into several distinct components, each responsible for specific aspects of the data collection and processing workflow, while maintaining loose coupling and high cohesion between components.
The core of the system architecture is built around the AIPerceptionTimeline class, which serves as the central orchestrator for the data collection process. This class implements the main data collection logic, managing the interaction with the Google Custom Search API and coordinating the collection of data across different time periods. The class is designed with a clear separation of concerns, implementing distinct methods for search query construction, result processing, and data persistence. This modular design allows for easy extension and modification of the system's capabilities as new requirements emerge.

The architecture includes a dedicated web_content_parsers module that implements specialized parsers for different content platforms and formats. This module follows a plugin-based architecture, where each parser is implemented as an independent component that can be easily added, removed, or modified without affecting the rest of the system. The parsers include specialized implementations for major platforms such as Reddit, LinkedIn, Quora, YouTube, and PDF documents, each handling the unique characteristics and challenges of their respective platforms. This modular approach to content parsing allows for flexible adaptation to different content sources and formats while maintaining consistent data quality and structure.

The system includes a comprehensive utils module that implements various utility functions and helper classes used throughout the system. This module includes specialized components for content translation, duplicate detection, data cleaning, and internet connectivity checking. The utility module is designed with a focus on reusability and maintainability, implementing common functionality that can be used across different parts of the system. This approach reduces code duplication and ensures consistent implementation of common operations.

The architecture implements a sophisticated data processing pipeline that handles the collected data through various stages of processing and analysis. The pipeline includes components for initial data collection, content extraction, data cleaning, and sentiment analysis. Each stage of the pipeline is implemented as a separate component with well-defined interfaces, allowing for easy modification or replacement of individual processing steps. This modular approach to data processing ensures that the system can be easily adapted to different analysis requirements or data formats.

The system includes a robust error handling and recovery mechanism that ensures reliable operation even in the face of various failure scenarios. The architecture implements comprehensive error detection and handling at each level of the system, from API interactions to data processing and storage. The error handling system includes sophisticated recovery strategies, such as automatic retry mechanisms, fallback options, and graceful degradation of functionality. This approach ensures that the system can continue operating effectively even when individual components encounter problems.

The architecture implements a sophisticated logging and monitoring system that provides detailed insights into the system's operation and performance. The logging system captures comprehensive information about various aspects of the system's operation, including API interactions, data processing steps, and error conditions. This information is crucial for debugging, performance optimization, and system maintenance. The monitoring system includes real-time progress tracking and status reporting, providing users with clear visibility into the system's operation.

The system includes a flexible data storage and management component that handles the persistence and organization of collected data. The architecture supports multiple data storage formats, including CSV and JSON, with automatic file management and versioning. The data management system includes sophisticated mechanisms for data validation, cleaning, and organization, ensuring that the stored data is consistent, complete, and easily accessible for analysis.

The architecture implements a sophisticated configuration management system that allows for flexible customization of the system's behavior. The configuration system includes support for various parameters such as API credentials, rate limits, and processing options, which can be easily modified to adapt the system to different requirements or environments. This flexibility is particularly valuable for adapting the system to different research needs or operational constraints.

The system architecture is designed to be scalable and maintainable, with clear separation of concerns and modular components. The architecture follows established software engineering principles such as the Single Responsibility Principle and the Open-Closed Principle, ensuring that the system can be easily extended and modified as new requirements emerge. This approach to system design ensures that the system remains maintainable and adaptable throughout its lifecycle.

The architecture includes comprehensive documentation and testing infrastructure that supports the development and maintenance of the system. The documentation includes detailed specifications of the system's components, interfaces, and behavior, making it easy for developers to understand and modify the system. The testing infrastructure includes unit tests, integration tests, and system tests that ensure the reliability and correctness of the system's operation.

### 3.3.2 Key Technologies

The implementation of the data collection and analysis system for understanding public perception of artificial intelligence is underpinned by a carefully selected suite of technologies, each chosen for its robustness, flexibility, and suitability for the specific challenges posed by large-scale, multi-source web data collection and processing. The system is architected as a hybrid solution, combining a powerful Python backend for automated, programmable data collection and processing, with a user-friendly Chrome extension frontend that enables interactive, browser-based data gathering.

**Python Backend Technologies**

At the heart of the backend lies the Python programming language, selected for its extensive ecosystem of libraries, ease of use, and strong support for data science and web automation tasks. The core data collection engine is built using Python 3.6+, leveraging the google-api-python-client library to interface with the Google Custom Search API. This library provides a reliable and well-documented means of programmatically querying Google Search, handling authentication, pagination, and error management in a way that is both scalable and maintainable.

For data manipulation and storage, the system employs the pandas library, which offers powerful data structures and functions for handling tabular data. This is essential for organizing search results, cleaning and transforming data, and exporting datasets in formats such as CSV and JSON. The use of pandas ensures that the data pipeline remains efficient and that large volumes of data can be processed with minimal overhead.

To provide a rich, interactive command-line experience, the backend utilizes the rich library, which enables the display of progress bars, tables, and status panels. This enhances usability, especially for long-running data collection tasks, by giving real-time feedback on progress, elapsed time, and any encountered issues.

The system's multi-platform content extraction capabilities are realized through a set of specialized libraries and APIs. For Reddit, the praw (Python Reddit API Wrapper) library is used to authenticate and fetch posts and comments, handling Reddit's complex thread structures and API rate limits. For Quora, LinkedIn, and general web content, the system uses a combination of requests for HTTP access and custom HTML/JSON parsing logic to extract relevant information. YouTube transcripts are retrieved using the youtube-transcript-api, which programmatically fetches video subtitles for further analysis. For academic and technical documents, the PyPDF2 library is employed to extract text from PDF files, handling multi-page documents and various encoding challenges.

The backend also incorporates advanced error handling, rate limiting, and retry logic, using libraries such as tenacity for robust API interaction and logging for comprehensive event tracking. This ensures that the system can gracefully handle network interruptions, API quota limits, and unexpected data formats without data loss or corruption.

For translation and language normalization, the backend can integrate with external translation APIs, ensuring that content in multiple languages can be processed and analyzed uniformly. Deduplication and data cleaning utilities are implemented as modular scripts within the utils directory, providing reusable tools for maintaining data quality.

The Chrome extension component of the system is built using standard web technologies—HTML5 for structure, CSS3 for styling, and JavaScript for interactivity. Its user interface allows users to specify search queries, year ranges, and other parameters, and then automates the process of collecting Google Search results directly within the browser. The extension parses search result pages, extracts relevant data fields, and provides real-time feedback through progress bars and logs. Users can pause, resume, and export the collected data as CSV files for further analysis. This browser-based approach is especially useful for interactive, language-specific, or region-specific data collection scenarios where API access may be limited, and it is designed to be accessible to users without programming expertise.

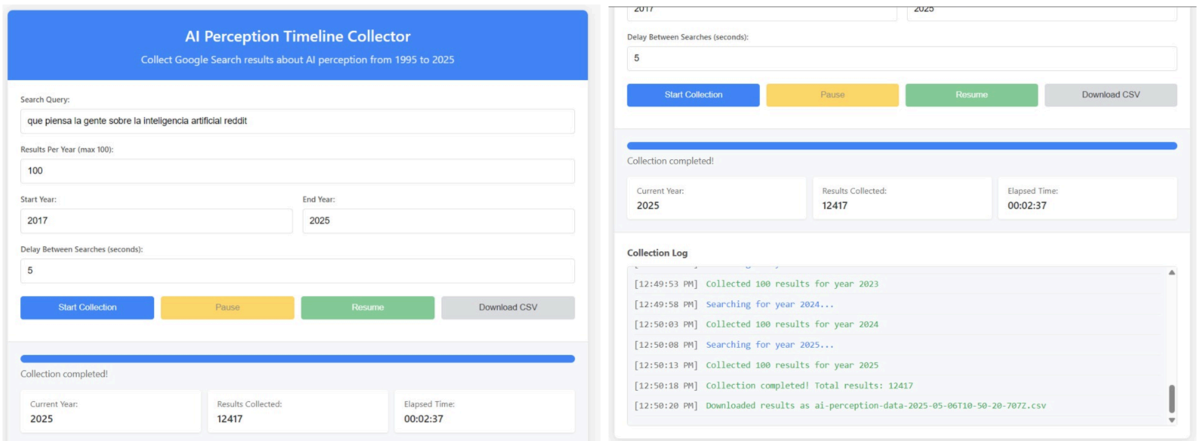


*Figure 1: Chrome Extension*

**Integration and Interoperability**

Both the backend and frontend components are designed to be interoperable, with data formats and structures standardized to facilitate seamless integration. Data collected via the Chrome extension can be exported and further processed by the Python backend, enabling advanced analysis, deduplication, and sentiment assessment. Conversely, the backend can generate datasets that are compatible with browser-based visualization and exploration tools.

The choice of these technologies ensures that the system is both powerful and flexible, capable of handling the scale and diversity of web data required for a comprehensive analysis of AI perception. The modularity of the architecture allows for easy extension, such as the addition of new content parsers or support for additional languages and platforms, ensuring that the system can evolve alongside the research needs.

## 3.4 Data Processing Pipeline

The data processing pipeline in this project is designed as a multi-stage workflow that transforms raw, heterogeneous web data into a structured, high-quality dataset suitable for advanced analysis, including sentiment and comparative studies. This pipeline ensures that the collected data is not only comprehensive but also clean, consistent, and analytically valuable.

The pipeline begins with the initial collection stage, where raw data is gathered from a variety of sources using both the Python backend and the Chrome extension. At this stage, the data is highly diverse, encompassing different formats, languages, and content structures. Each record typically includes metadata such as the year of publication, source domain, URL, title, and a snippet or main content body. The system ensures that every entry is tagged with its temporal and source context, which is crucial for subsequent longitudinal and cross-sectional analyses.

```
year_query = f"{query} after:{year-1} before:{year+1}"

for i in range(num_calls):
    start_index = i * 10 + 1
    if start_index > max_results:
        break

    try:
        response = (
            self.service.cse()
            .list(
                q=year_query,
                cx=self.cx,
                start=start_index,
            )
            .execute()
        )

        if "items" in response:
            for item in response["items"]:
                if len(results) >= max_results:
                    break

                # Extract domain from URL
                domain = item.get("displayLink", "")

                result = {
                    "year": year,
                    "title": item.get("title", ""),
                    "domain": domain,
                    "url": item.get("link", ""),
```

```
            "snippet": item.get("snippet", ""),
        }
        results.append(result)

    time.sleep(1)  # Wait between API calls
```

Following collection, the content processing stage is initiated. Here, the pipeline applies a series of normalization and cleaning procedures to the raw data. This includes the removal of HTML tags, special characters, and extraneous whitespace, as well as the unification of text encoding to ensure compatibility across different languages and platforms. For multilingual content, the system integrates translation utilities, enabling the normalization of all text into a common language (typically English) to facilitate uniform sentiment analysis. The pipeline also incorporates language detection algorithms to accurately identify and process content in various languages, ensuring that no relevant data is excluded due to linguistic diversity.

A critical component of the processing stage is the deduplication mechanism. Given the scale and overlap of web content, it is common to encounter duplicate or near-duplicate articles, especially when aggregating from multiple sources. The pipeline employs both exact and fuzzy matching algorithms to identify and remove redundant entries, relying on a combination of content similarity, URL matching, and metadata comparison. This step is essential for maintaining the integrity of the dataset and preventing bias in the analytical results.

Once the data is cleaned and deduplicated, the pipeline proceeds to the quality control phase. This involves a series of validation checks to ensure that each record meets predefined standards for completeness and relevance. The system verifies the presence of essential fields such as year, content, and source, and flags or removes entries that are empty, malformed, or irrelevant to the research focus. Additional checks are performed to confirm the accuracy of language detection, the success of translation, and the consistency of metadata.

The final output of the data processing pipeline is a structured dataset, typically in CSV or JSON format, that is ready for advanced analytical procedures such as sentiment analysis, time series analysis, and demographic comparisons. The pipeline is designed to be modular and extensible, allowing for the integration of additional processing steps or analytical modules as research needs evolve. Throughout the entire process, comprehensive logging and error tracking are maintained, ensuring transparency, reproducibility, and ease of troubleshooting.

## 3.5 Ethical Considerations

The ethical dimension of this research is addressed through careful attention to data privacy, responsible use of technology, and compliance with the terms of service of all platforms involved. Throughout the data collection process, only publicly available information is accessed, and no attempts are made to bypass paywalls, authentication barriers, or privacy settings. The use of official APIs, such as Google Custom Search and Reddit's PRAW, ensures that data is gathered in accordance with the platforms' intended usage policies and rate limits, minimizing the risk of service disruption or unauthorized access.

To further protect privacy, the system is designed to avoid collecting or storing personally identifiable information (PII) beyond what is already publicly visible in the context of web articles or forum posts. Any incidental PII that appears in the collected data is handled with care, and, if necessary, can be redacted or excluded during the data cleaning phase. The project also respects copyright and intellectual property rights by using content solely for academic research and analysis, and by providing proper attribution to original sources in any published results.

Finally, the research acknowledges the potential for bias in web data and the ethical implications of automated sentiment analysis. All findings are interpreted with an awareness of these limitations, and the methodology is transparently documented to allow for critical evaluation and reproducibility.

## 3.6 Methodological Limitations

While the data collection and analysis framework developed for this research is robust and versatile, several methodological limitations must be acknowledged. The reliance on public APIs and web scraping means that the dataset is inherently constrained by the availability and accessibility of online content, as well as by the rate limits and coverage of the APIs used. Some platforms may restrict access to certain types of content or limit the number of results returned, which can introduce sampling bias or gaps in the data, particularly for earlier years or less-represented languages and regions.

Additionally, the process of deduplication and content cleaning, while thorough, may not perfectly eliminate all redundant or irrelevant entries, especially when articles are republished across multiple domains with slight modifications. Language detection and automated translation, though highly effective, are not infallible and may introduce errors or subtle shifts in meaning that could affect sentiment analysis outcomes.

## 3.7 Data Validation and Quality Assurance

Ensuring the reliability and integrity of the collected dataset is a central priority throughout the research process. Data validation is performed at multiple stages, beginning with the initial collection, where each record is checked for completeness, correct metadata, and relevance to the research topic. Automated scripts verify that essential fields such as year, source, and content are present and that URLs are valid and accessible.

During the cleaning and deduplication phases, both exact and fuzzy matching techniques are used to identify and remove duplicate or near-duplicate entries, reducing redundancy and potential bias. Language detection and translation steps are validated by cross-checking detected languages and reviewing a sample of translated content for accuracy. The sentiment analysis pipeline incorporates checkpoints and logging to monitor for processing errors or anomalies, allowing for rapid identification and correction of issues.

Quality assurance is further supported by periodic manual reviews of random data samples, ensuring that automated processes are functioning as intended and that the dataset remains representative and free from systematic errors. All validation steps and quality checks are documented, providing transparency and reproducibility for future research or audits.

# 4. Results

## 4.1 Overview of the Dataset

The dataset underpinning this analysis comprises a curated collection of 4,091 online articles, blog posts, and commentary pieces spanning from 2018 to 2025, each addressing various dimensions of artificial intelligence. This corpus, drawn from a wide array of web domains, includes content originally published in multiple languages, subsequently standardized for sentiment analysis and comparative study. Each entry in the dataset contains a rich set of metadata, including the year of publication, source domain, and URL, alongside the article title and a brief snippet, which together offer contextual grounding and temporal traceability of AI discourse online.

To ensure analytical depth, the dataset includes fields that denote the source type (e.g., web article, LinkedIn post), the inferred gender of the author or speaker, and the language of the original content. Notably, articles not originally written in English were machine-translated and stored in a dedicated translatedContent field, with a Boolean isTranslated flag facilitating distinction between native and translated texts. The cleaned_content and clean_ai_focus fields contain processed versions of the full article texts, stripped of extraneous formatting or noise, tailored specifically for natural language processing tasks such as sentiment and aspect-based analyses.

Each text is annotated with multiple layers of sentiment information. The overall_sentiment field classifies the entire piece as positive, negative, or neutral, whereas aspect_sentiments and aspects provide a granular breakdown of sentiment polarity across thematic dimensions such as “technology,” “bias,” “access,” and “human relationships.” This design enables both holistic and focused interpretations of public attitudes toward AI. Further emotional categorization is encoded under the emotion variable, allowing for a cross-sectional look into affective reactions such as fear, confusion, or enthusiasm associated with the topic. A synthesized summary of each text is also available, facilitating high-level comprehension and aiding in qualitative validations of the automated sentiment assessments.

This dataset retains a representative sample of the linguistic, temporal, and demographic diversity embedded in the broader collection. It includes entries in English, Italian, Spanish and French, and reflects varying perspectives from different gender. One field, clean_ai_focus, is missing a single entry, indicating an almost complete cleaning and preprocessing process.

In sum, this dataset presents a nuanced and multidimensional resource for evaluating public discourse on artificial intelligence, capturing not only the thematic breadth of discussions but also the evolution of sentiment and emotion across different contexts and over time.

## 4.2 Temporal Trends in AI Sentiment

The temporal evolution of sentiment toward artificial intelligence was analyzed across the six-year period from 2018 to 2025, with particular attention to both overall trends and gender-based differences.

Figure 2 presents the yearly average sentiment scores for articles authored by males and females. The results reveal that, while both groups generally fluctuate around a neutral stance, male-authored articles consistently exhibit slightly higher average sentiment scores compared to female-authored ones. Notably, sentiment among male authors shows a modest upward trend from 2020 to 2022, peaking in 2022, before experiencing a mild decline in the subsequent years. Female-authored articles, in contrast, display greater volatility, with sentiment dipping below neutral in 2021 and 2023, suggesting episodes of increased skepticism or concern within this demographic. Despite these fluctuations, the overall gap between male and female sentiment remains relatively modest throughout the observed period.

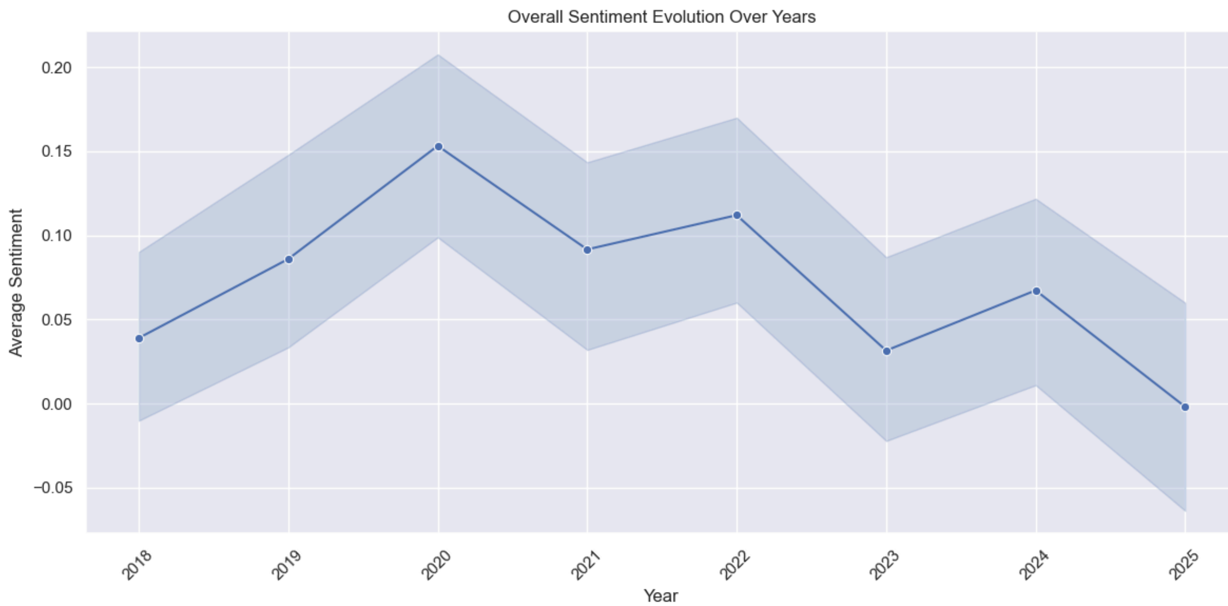


*Figure 2: Overall sentiment Evolution Over Years*

Figure 3 further dissects these dynamics by illustrating the average sentiment for each gender in every year. The bar plot highlights that the most pronounced gender differences occur in 2021 and 2023, where female sentiment is notably more negative than male sentiment. In other years, the averages for both genders converge closer to the overall mean, indicating that the observed divergences are not persistent but rather episodic, potentially linked to specific events or shifts in public discourse during those years.

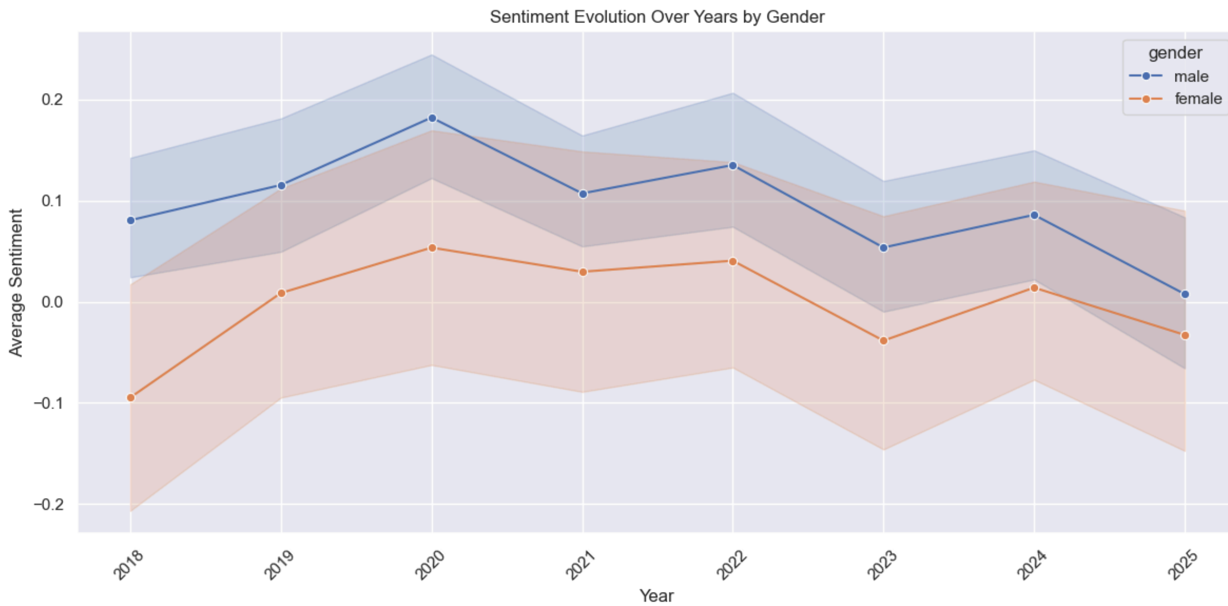


*Figure 3: Sentiment evolution over years by gender*

Turning to the aggregate perspective, Figure 4 depicts the overall average sentiment toward AI for each year, irrespective of gender. The general trend suggests a gradual increase in positive sentiment from 2019 through 2022, followed by a slight downturn in 2023 and 2024. This trajectory may reflect an initial period of optimism or enthusiasm about AI developments, tempered in later years by growing public debate, regulatory scrutiny, or high-profile controversies.

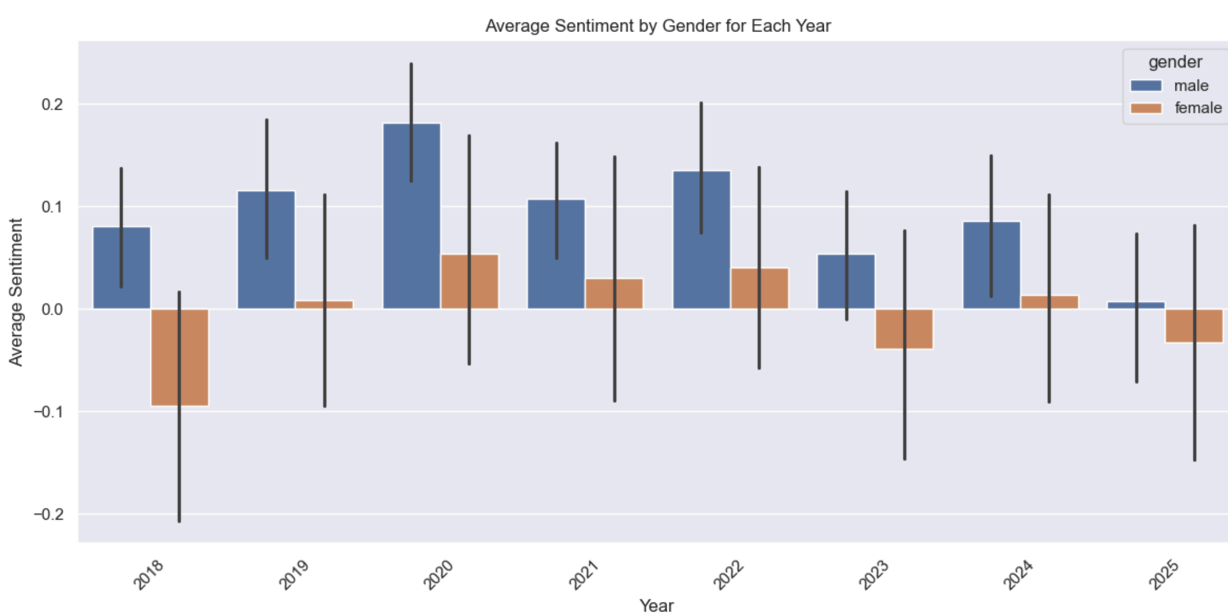

*Figure 4: Average sentiment by gender for each year*

Collectively, these visualizations indicate that while sentiment toward artificial intelligence in online discourse remains close to neutral on average, there are discernible temporal patterns and gender-based nuances. The data suggest periods of heightened optimism as well as intervals marked by increased caution or negativity, with gender differences most evident during years of greater sentiment volatility.

## 4.3 Demographic and Linguistic Variation

This section explores how sentiment toward artificial intelligence varies across different demographic and linguistic groups, as revealed by several key visualizations.

Figure 5 presents the average sentiment associated with each language in the dataset. This bar plot offers a comparative view of how sentiment toward artificial intelligence differs across linguistic communities, highlighting potential cultural or regional distinctions in public perception.

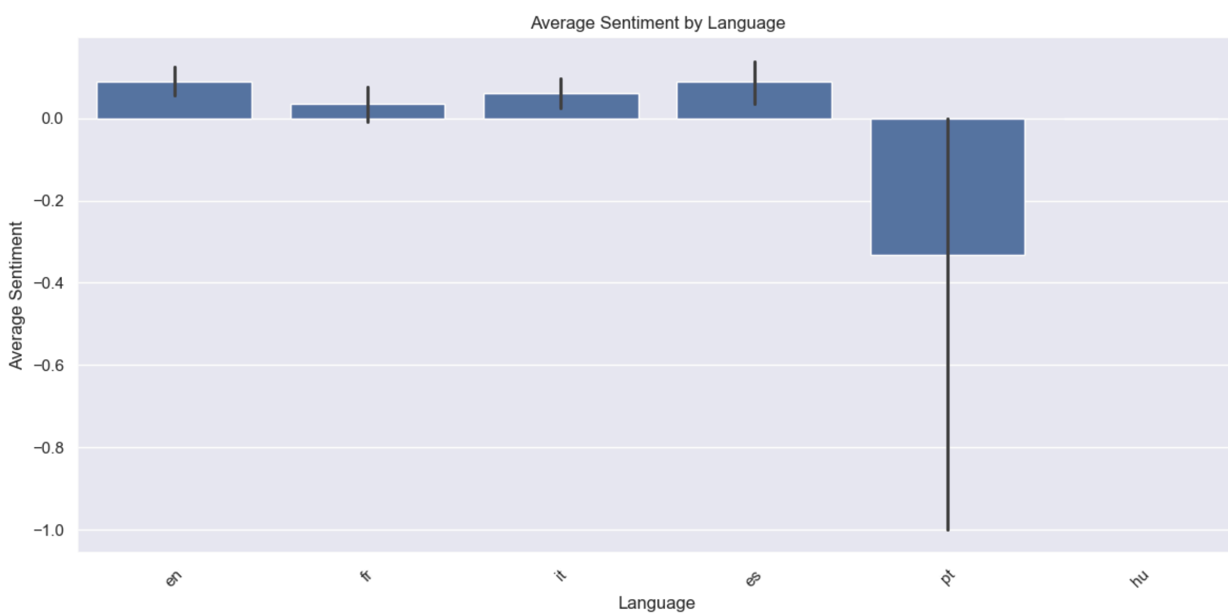


*Figure 5: Average sentiment by language*

*Figure 6 displays the evolution of sentiment over time for each language group. By tracking sentiment trends within individual language communities, this line plot allows for the observation of both stable and shifting attitudes toward AI, and can reveal whether certain languages exhibit consistent perspectives or respond dynamically to broader developments in the field.*

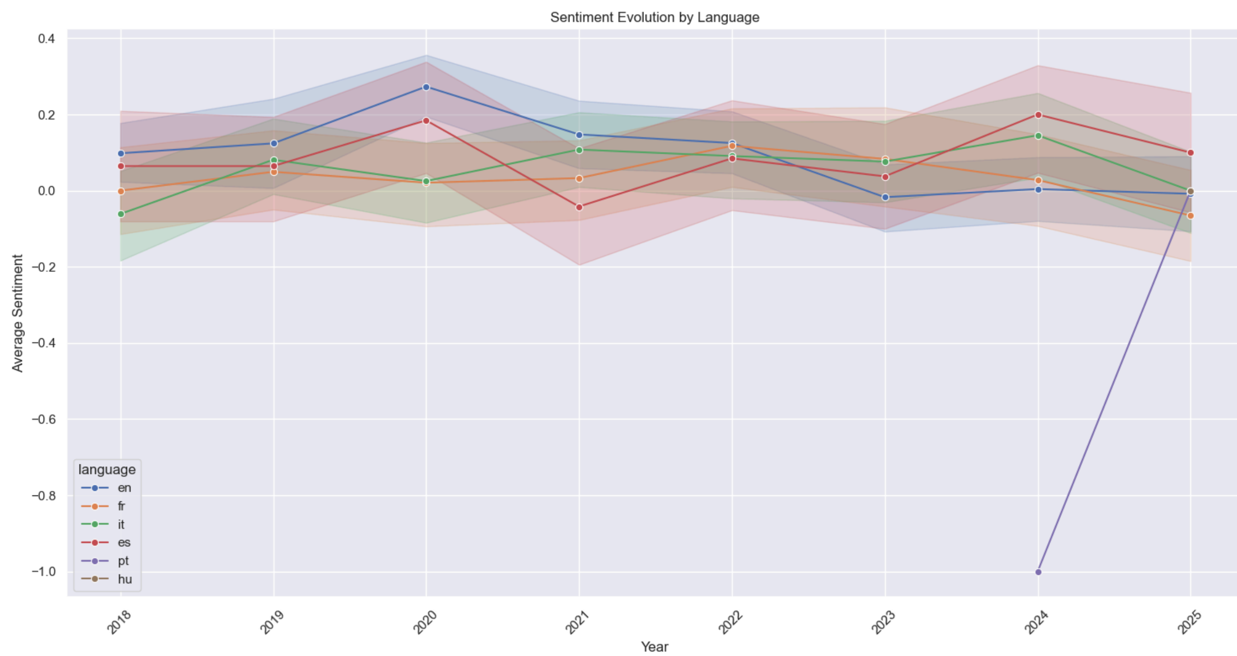


*Figure 6: Sentiment evolution by language*

Figure 7.1 and 7.2 show the distribution of sentiment classes—positive, neutral, and negative—by year and gender, using stacked bar plots. This visualization enables a clear comparison of how the proportions of sentiment categories vary annually within each gender group, thereby providing insight into demographic patterns and changes in public attitudes over time.

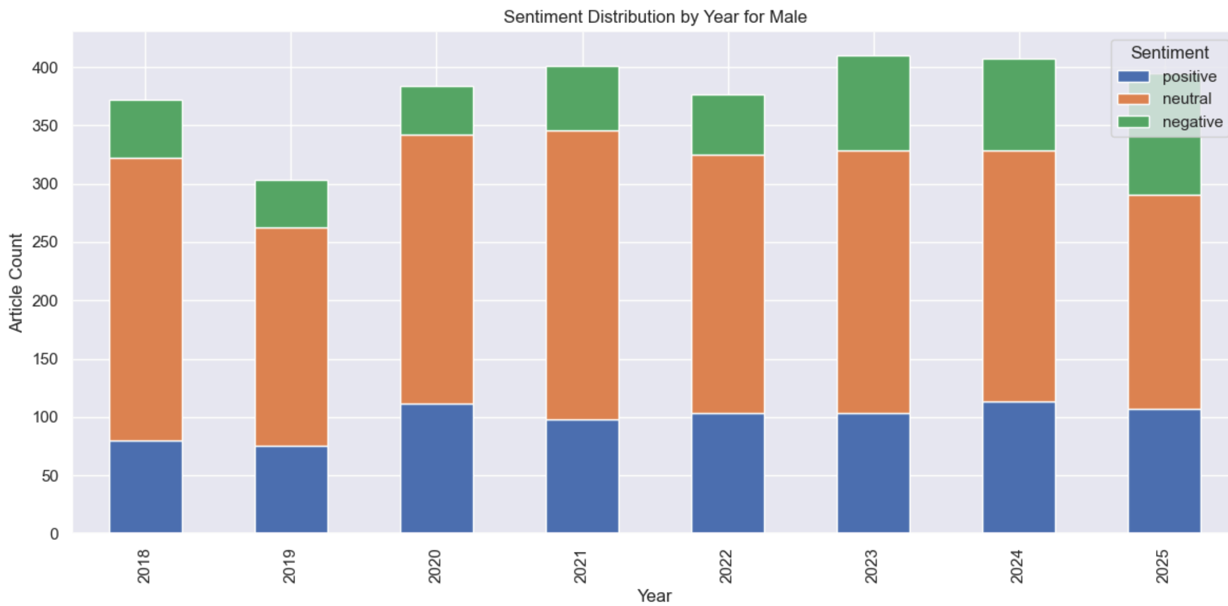


*Figure 7.1: Sentiment distribution by year and for male*

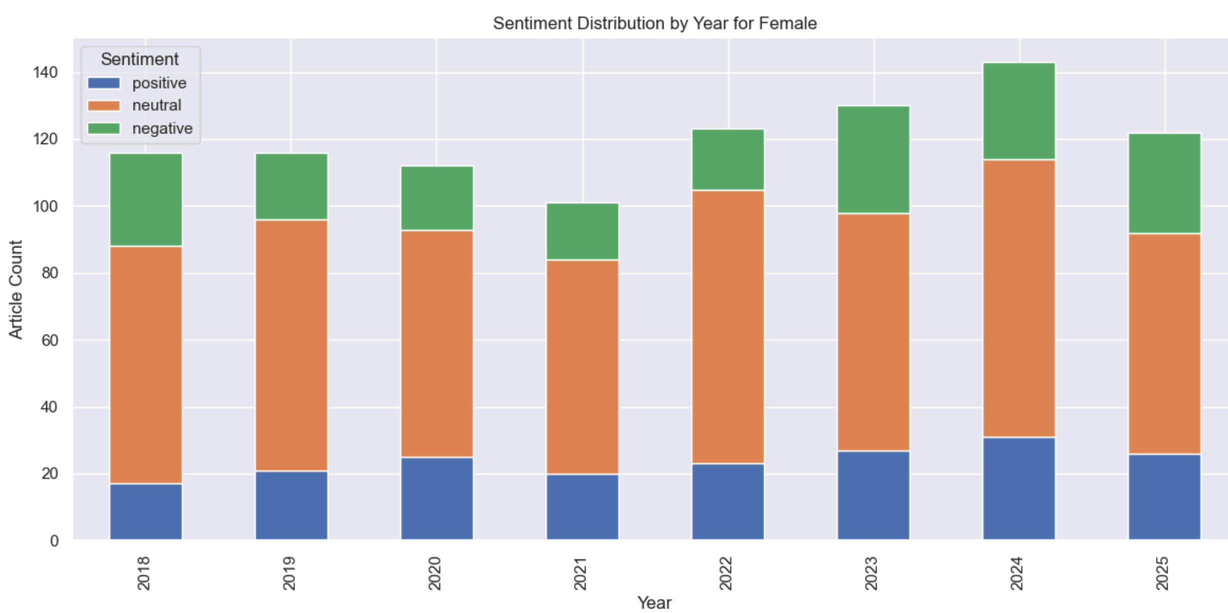


*Figure 7.2: Sentiment distribution by year and for female*

Figure 8 illustrates the sentiment class distribution by language, again using a stacked bar plot. This figure helps clarify whether particular languages are characterized by a predominance of a specific sentiment class, or whether sentiment is more evenly distributed, thus offering a nuanced perspective on the linguistic dimension of AI perception.

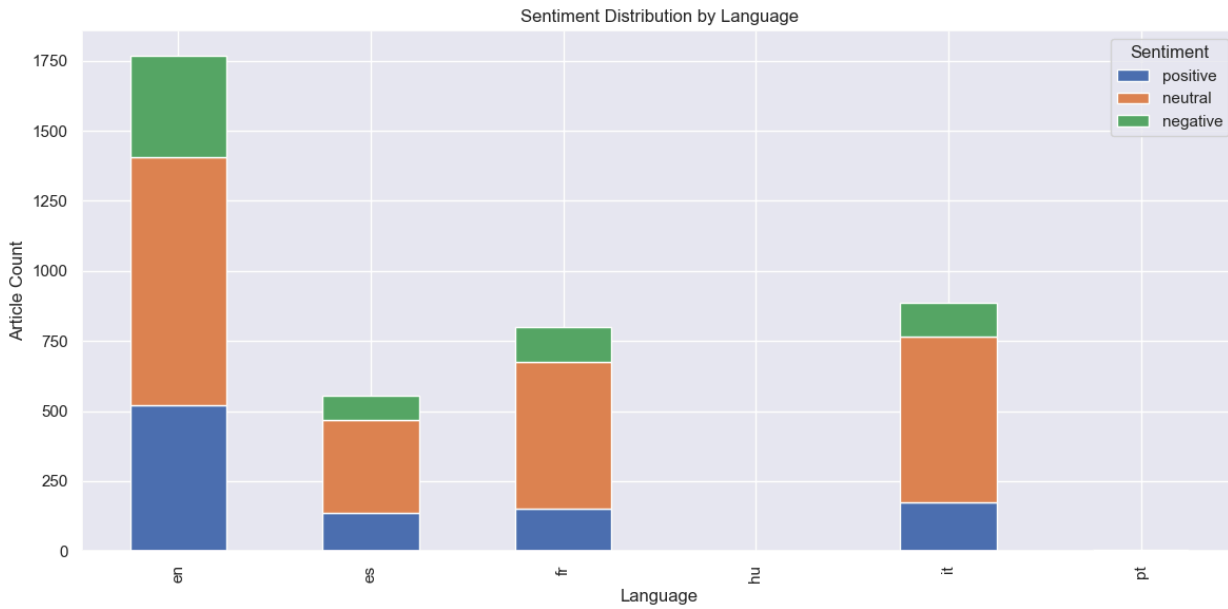


*Figure 8: Sentiment distribution by language*

Taken together, Figures 5 through 8 demonstrate that sentiment toward artificial intelligence is not uniform across demographic and linguistic groups. The visualizations reveal clear differences in both the average sentiment and the distribution of sentiment classes when the data is segmented by language and gender. These patterns suggest that cultural, regional, and demographic factors play a significant role in shaping public attitudes toward AI. The observed variations highlight the importance of considering such factors in any comprehensive analysis of AI perception, and they provide a foundation for more nuanced interpretation and discussion in subsequent sections.

## 4.4 Emotional and Thematic Insights

Figure 9 presents the dominant emotion trends over time within the dataset. This line plot tracks the prevalence of various emotions—such as hope, fear, or enthusiasm—expressed in AI-related articles across the years. By visualizing how the frequency of specific emotions changes over time, this figure provides a high-level view of the evolving emotional landscape in public discourse on artificial intelligence. To offer a more granular perspective, additional figures that break down these emotion trends by gender and by language are included in the appendix.

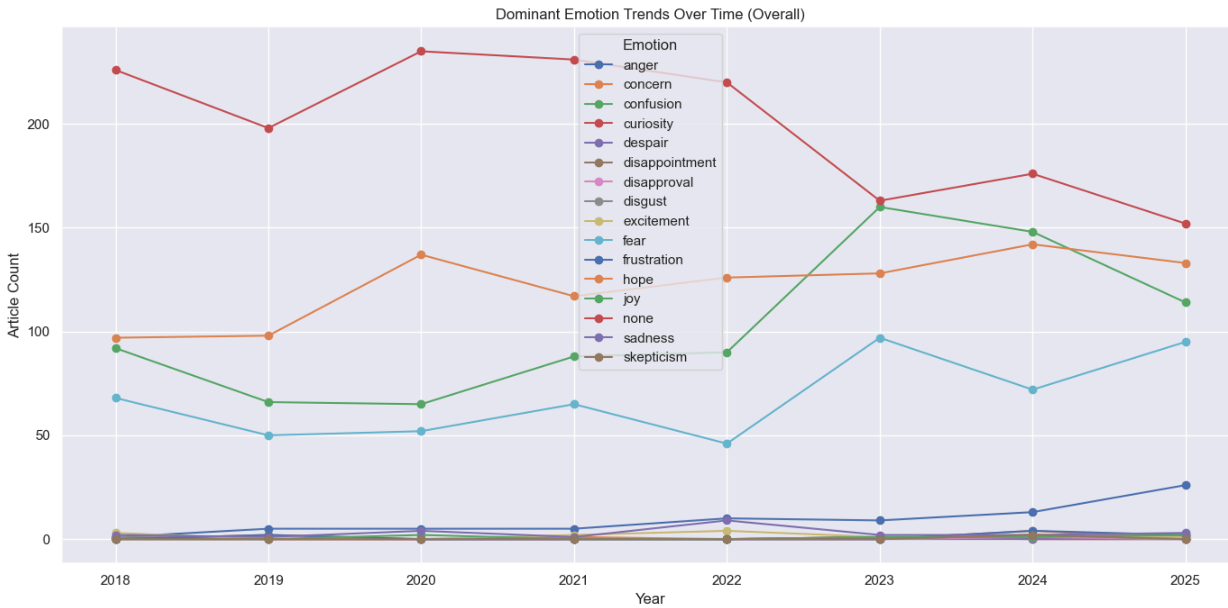


Figure 9: Dominant emotion trends over time

Figure 10 illustrates the results of the aspect-based sentiment analysis, which examines sentiment polarity associated with the most frequently discussed AI-related aspects or themes. This bar plot highlights which topics—such as ethics, innovation, or job impact—are generally perceived positively or negatively within the corpus. By focusing on sentiment at the aspect level, this visualization uncovers which dimensions of AI are sources of optimism or concern in public discussions.

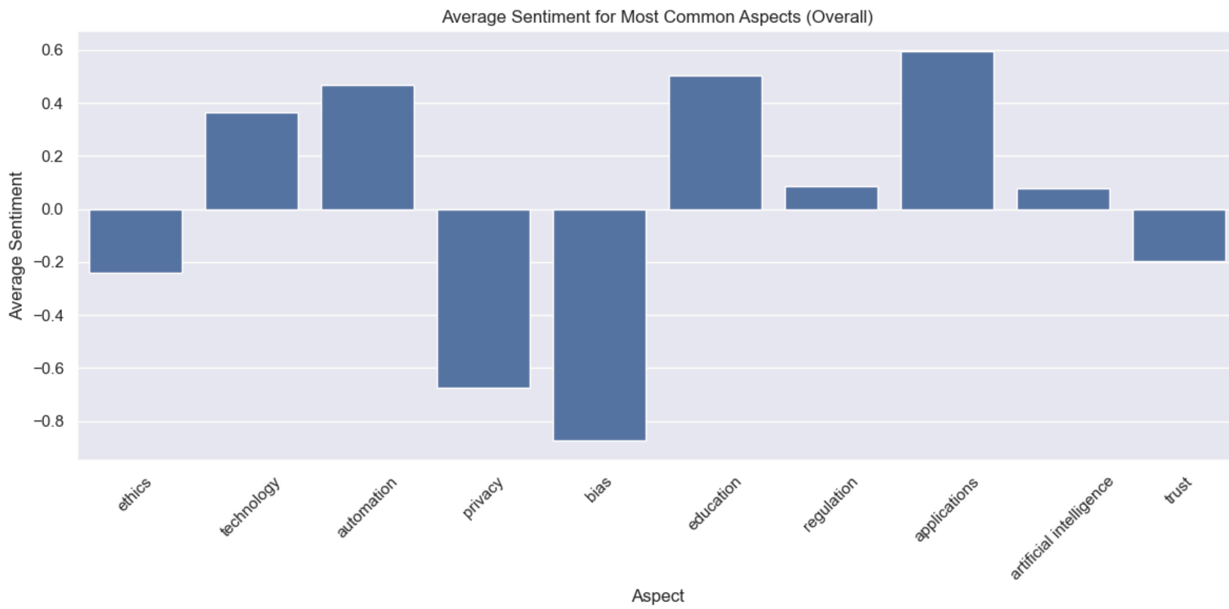


Figure 10: Aspect-based sentiment analysis

*In summary, Figures 9 and 10 together provide insight into not only the overall emotional tone of AI discourse but also the specific themes that drive positive or negative sentiment. The inclusion of more detailed emotion trend analyses in the appendix allows for a deeper exploration of how these patterns differ across demographic and linguistic groups, further enriching the understanding of public perceptions surrounding artificial intelligence.*

## 4.5 Overall Sentiment Distribution

Figure 11 presents the overall distribution of sentiment classes—positive, neutral, and negative—across the entire dataset, visualized as a pie chart. This figure provides a high-level summary of the general mood expressed in public discourse about artificial intelligence during the study period. By illustrating the proportion of articles falling into each sentiment category, the visualization offers an immediate sense of whether the prevailing tone in the collected content is optimistic, balanced, or marked by skepticism.

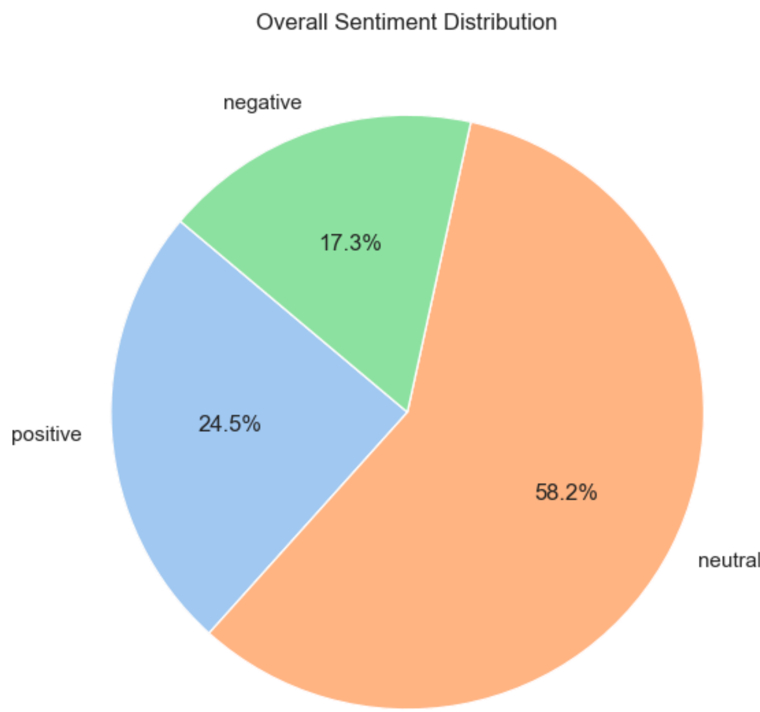


Figure 11: Overall sentiment distribution

This summary perspective serves as a useful reference point for interpreting the more detailed analyses presented in earlier sections. It helps contextualize the findings on temporal, demographic, and thematic variation by showing the overall sentiment landscape in which these patterns emerge.

## Summary of Findings

The analyses presented in this chapter offer a multidimensional portrait of how artificial intelligence is perceived in online discourse, drawing on a diverse and richly annotated dataset of articles, posts, and commentaries spanning several years, languages, and demographic groups.

Temporal analyses reveal that sentiment toward AI has not remained static over the observed period. Instead, there are discernible fluctuations, with certain years marked by increased optimism or heightened skepticism, often corresponding to broader developments in AI technology or public debate. While the overall sentiment tends to hover near neutrality, periods of both heightened positivity and negativity are evident, suggesting that public attitudes are responsive to external events and evolving narratives within the field.

Demographic and linguistic analyses underscore the importance of cultural and social context in shaping AI perception. Sentiment varies not only by gender—with some differences in the degree and direction of sentiment between male and female authors—but also by language, indicating that regional or cultural factors influence how AI is discussed and evaluated. The distribution of sentiment classes further highlights that some language groups or demographic segments are more likely to express positive, neutral, or negative views, reinforcing the need to consider these variables in any comprehensive assessment of public opinion.

Thematic and emotional analyses add further nuance to these findings. The tracking of dominant emotions over time reveals that the emotional tone of AI discourse is dynamic, with certain emotions—such as hope, fear, or enthusiasm—rising or falling in prominence depending on the year and, as detailed in the appendix, the demographic or linguistic subgroup. Aspect-based sentiment analysis demonstrates that public attitudes are not monolithic; instead, sentiment varies considerably across different AI-related themes. For example, topics such as innovation and opportunity may be discussed in a more positive light, while issues related to ethics, bias, or societal impact can evoke more critical or concerned sentiment.

Finally, the overall sentiment distribution provides a high-level summary of the dataset, showing the relative proportions of positive, neutral, and negative sentiment present in the collected material. This aggregate perspective contextualizes the more granular findings, confirming that while there is a balance among sentiment classes, neither uncritical optimism nor pervasive negativity dominates the discourse.

Taken together, these results highlight the complexity and diversity of public perception regarding artificial intelligence. They demonstrate that sentiment is shaped by a confluence of temporal, demographic, linguistic, and thematic factors, and that understanding these patterns requires both broad quantitative analysis and attention to the underlying context. These findings lay the groundwork for a deeper discussion of their implications, limitations, and relevance to ongoing debates about the societal impact of AI.

# 5. Conclusion

## 5.1 Summary of Key Findings

This thesis presented a multidimensional analysis of public attitudes toward artificial intelligence, drawing on a large-scale, multilingual dataset from 2018 to 2025 consisting of multiple types of sources, languages, and demographic groups. The findings showed that sentiment toward AI is dynamic and context-specific, bounded by technological events, cultural context, and demographic factors. Overall, tone around AI discussions seems to stay close to neutral, but there are periods of optimism and skepticism, often timed with major events, or policy discussions. Gender and language are major moderators of sentiment, with both groups demonstrating notable differences between the average sentiment and the distribution of sentiment classes. Thematic, and emotional analyses indicate that public attitudes toward AI are not uniform, but shaped by particular concerns – ethics, jobs, and innovation, as well as emotional tones of hope, fear, and uncertainty that feature prominently in AI discussions. The study also emphasized the role of source diversity and translation, and the complicated ways in which policy, language, and opinion intersect.

## 5.2 Contributions to the Field

This study makes considerable contributions to the literature on AI perception in several important ways. First, it demonstrates the value of a multi-source, multi-language study that captures nuances and breadth of public sentiment, which are often limited by single-platform or single-language studies. By leveraging web articles, social media, professional network data, and academic contents, this study provides a fuller picture of how AI discourse evolves over multiple representations across different communities. Second, the study demonstrates that advanced sentiment analysis techniques (including aspect-based sentiment and emotion methods) allow for a more nuanced understanding of both theme and emotional response to public sentiment. Third, the methodological rigor of this approach, particularly with regards to data cleaning, deduplication, and validation, serves as a replicable investigation design and contributes to large-scale sentiment studies. Finally, this study focuses on demographic and cultural variation and has implications for the wider discourse about the global and local characteristics of AI acceptance and skepticism.

## 5.3 Limitations of the Study

Despite its strengths, this study is subject to several limitations. The dataset, while extensive, is not exhaustive and may be affected by selection bias, particularly in the coverage of non-English sources and underrepresented regions. Machine translation, though carefully validated, may introduce subtle distortions in sentiment or meaning, especially for idiomatic or culturally specific expressions. The reliance on automated gender inference and source metadata may also result in occasional misclassification. Furthermore, sentiment analysis models, even those employing advanced NLP techniques, can struggle with ambiguity, sarcasm, and mixed sentiment, potentially leading to misinterpretation of nuanced opinions. Finally, the cross-sectional nature of much of the data limits the ability to draw causal inferences about the drivers of sentiment change over time.

## 5.4 Recommendations for Future Research

Future studies should aim to address these limitations by expanding the linguistic and regional diversity of datasets, incorporating more native-language sources from the Global South and other underrepresented contexts. The adoption of longitudinal designs would enable a deeper understanding of how attitudes toward AI evolve in response to technological, political, or societal shifts. Methodological improvements in sentiment analysis—such as the integration of context-aware and culturally sensitive models, as well as the use of human-in-the-loop validation—would enhance the accuracy and interpretability of findings. Finally, as AI systems themselves become more capable of generating and shaping public discourse, critical attention should be paid to the feedback loops between AI, media, and societal attitudes.

## 5.5 Final Remarks

The fast adoption of artificial intelligence technologies brings unprecedented opportunities and challenges for societies worldwide. Understanding how people perceive AI—across languages, cultures, and communities—is vital so that we can innovate responsibly and earn public trust when doing so. This thesis highlights the complexity of perception of AI, revealing a landscape characterized by optimism and caution, driven by different voices and to an extent in a changing time. By highlighting the patterns and drivers of sentiment in AI-related conversations across the globe, this thesis contributes to an informed and comprehensive conversation about the technology and its role in society. As AI continues to evolve, it is essential that we keep learning and talking about its development so it meets the values, needs, and hopes of all involved.

References


1. Cognizant. (2024). Gen AI's Impact on Women in the Workplace. https://www.cognizant.com/us/en/insights/insights-blog/gen-ai-impact-on-women-in-the-workplace-wf2458851
2. Do Better Esade. (2024). AI and Cultural Differences. https://dobetter.esade.edu/en/AI-cultural-differences
3. Frontiers in Psychology. (2023). A criticality map to understand bias in the public perception of AI. https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2023.1190326/full
4. Future of Life Institute. (2023). Op-Ed: Poll Shows Strong Support for AI Regulation. https://futureoflife.org/recent-news/op-ed-poll-shows-strong-support-ai-regulation/
5. HAI Stanford. (2024). How Culture Shapes What People Want from AI. https://hai.stanford.edu/news/how-culture-shapes-what-people-want-ai
6. HAI Stanford. (2025). AI Index Report 2025, Chapter 8. https://hai-production.s3.amazonaws.com/files/hai_ai-index-report-2025_chapter8_final.pdf
7. Insight7. (2025). Aspect-Based Sentiment Analysis: A Detailed Explanation. https://insight7.io/aspect-based-sentiment-analysis-detailed-explanation/
8. Insight7. (2025). Challenges in Sentiment Analysis: Key Considerations. https://insight7.io/challenges-in-sentiment-analysis-key-considerations/
9. John Whelchel. (2024). Sentiment Analysis of Foreign Documents to Detect Translation Bias. https://johnwhelchel.com/static/mysite/index/IWBiases.pdf
10. McKnight's. (2023). Survey: Older Adults Don't Trust AI for Their Health Information. https://www.mcknights.com/news/survey-older-adults-dont-trust-ai-for-their-health-information/
11. PeerJ Computer Science. (2025). Kumar, M., Khan, L. & Chang, H.-T. Evolving techniques in sentiment analysis: a comprehensive review. https://peerj.com/articles/cs-2592.pdf
12. Penfriend AI. (2024). Sentiment Analysis: The Complete Guide. https://penfriend.ai/blog/sentiment-analysis

13. Pew Research Center. (2023). Americans' Views of Artificial Intelligence. https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/
14. Pew Research Center. (2023). Public Awareness of Artificial Intelligence in Everyday Activities. https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/
15. Pew Research Center. (2025). How the US Public and AI Experts View Artificial Intelligence. https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/
16. Software Oasis. (2024). AI Public Perception. https://softwareoasis.com/ai-public-perception/
17. arXiv. (2024). AI Perceptions Across Cultures: Similarities and Differences in Expectations, Risks, Benefits, Tradeoffs, and Value in Germany and China. https://arxiv.org/html/2412.01459v1

# Appendix

Figure 12 and 13 present the dominant emotion trends over time for male and female authors in AI-related discourse. These visualizations reveal both commonalities and differences in the emotional landscape associated with each gender. For male-authored articles, emotions such as "hope" and "fear" consistently appear as the most prevalent across the years, with "hope" and "fear" often fluctuating in response to broader developments in AI. In contrast, female-authored articles display a more varied emotional profile, with "fear," "hope," and "excitement" emerging as prominent emotions, and noticeable shifts in their relative frequency over time. Both groups show low but persistent levels of negative emotions such as "anger," "sadness," and "disappointment," though these are generally less frequent than neutral or positive emotions. The comparative analysis suggests that while hope and fear are central to AI discourse for all genders, female-authored content may exhibit greater emotional diversity and responsiveness to specific events or themes. These patterns underscore the importance of considering gender as a moderating factor in the affective dimensions of public attitudes toward artificial intelligence.

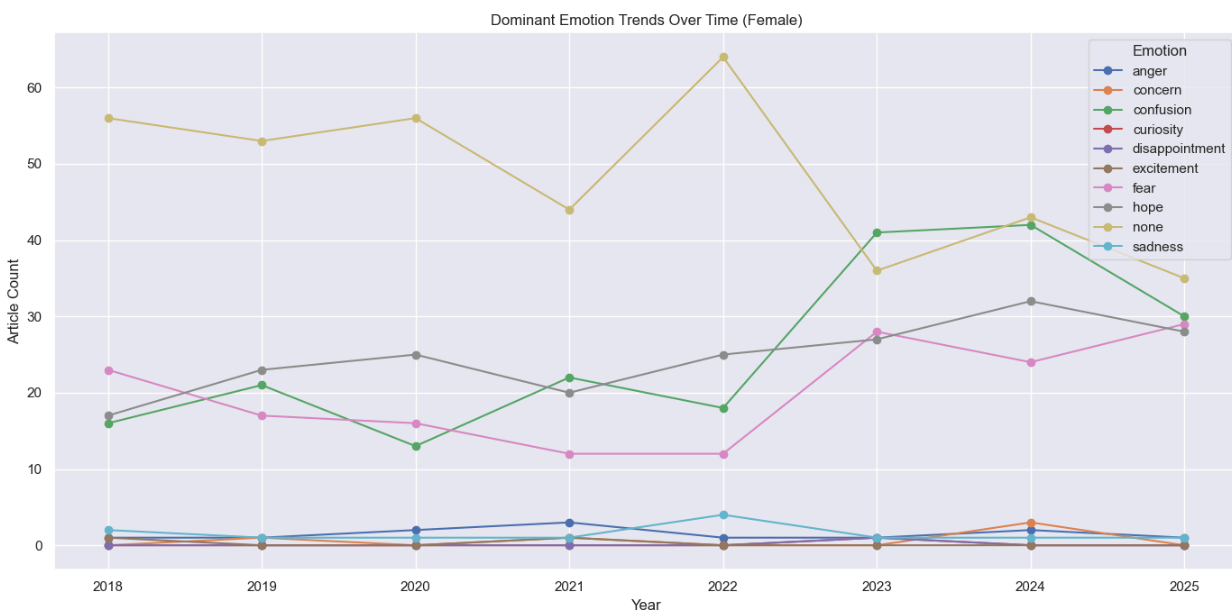


*Figure 12: Dominant Emotions Trends Over Time (Female)*

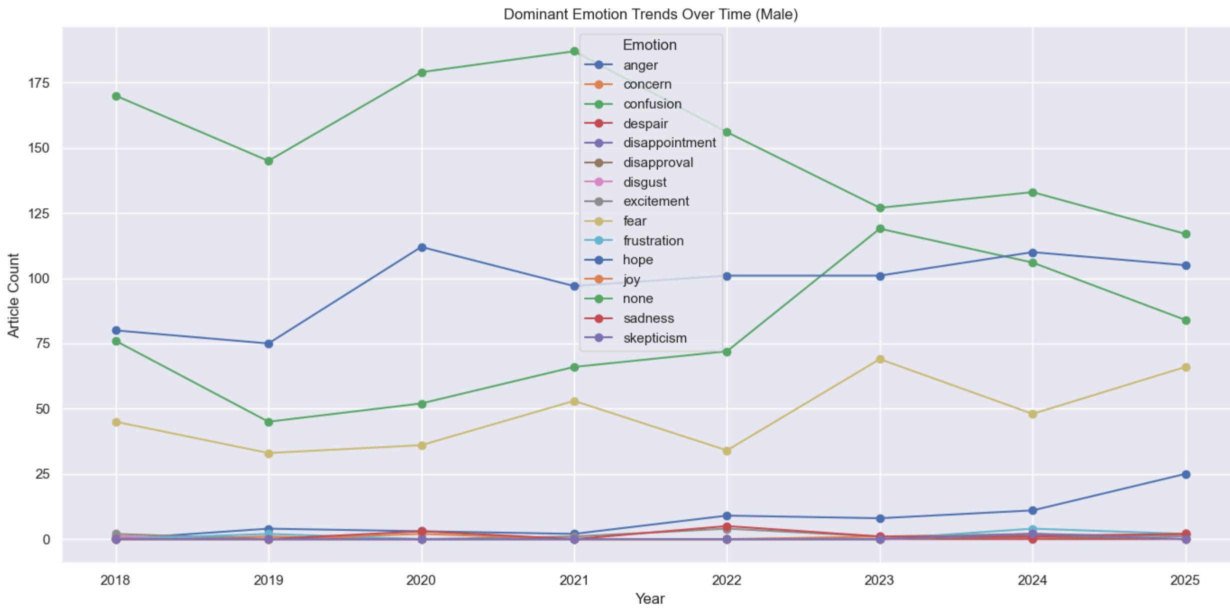


*Figure 13: Dominant Emotions Trends Over Time (Male)*

The following figures (figures 14-17) include a series of visualizations depicting dominant emotion trends over time for English, French, Italian, and Spanish language articles in the dataset. In English-language content, “hope” and “fear” are consistently among the most prevalent emotions, with their relative frequency shifting in response to broader developments in AI, while “confusion” and “disapproval” also appear at various points. French-language articles show a persistent presence of “confusion,” and “hope,” with “fear” and “disappointment” occasionally rising in prominence, suggesting a more varied but generally moderate emotional landscape. In Italian articles, “hope” and “excitement” are also significant, and there is a visible increase in “sadness” and “fear” in later years, indicating a shift toward more emotionally charged discourse. Spanish-language content displays a more dynamic pattern, with “hope,” and “fear” alternating as leading emotions, and occasional spikes in “confusion” and “frustration.” Overall, these trends highlight both shared and language-specific emotional responses to AI, reflecting cultural and contextual influences on how AI is discussed and perceived in different linguistic communities.

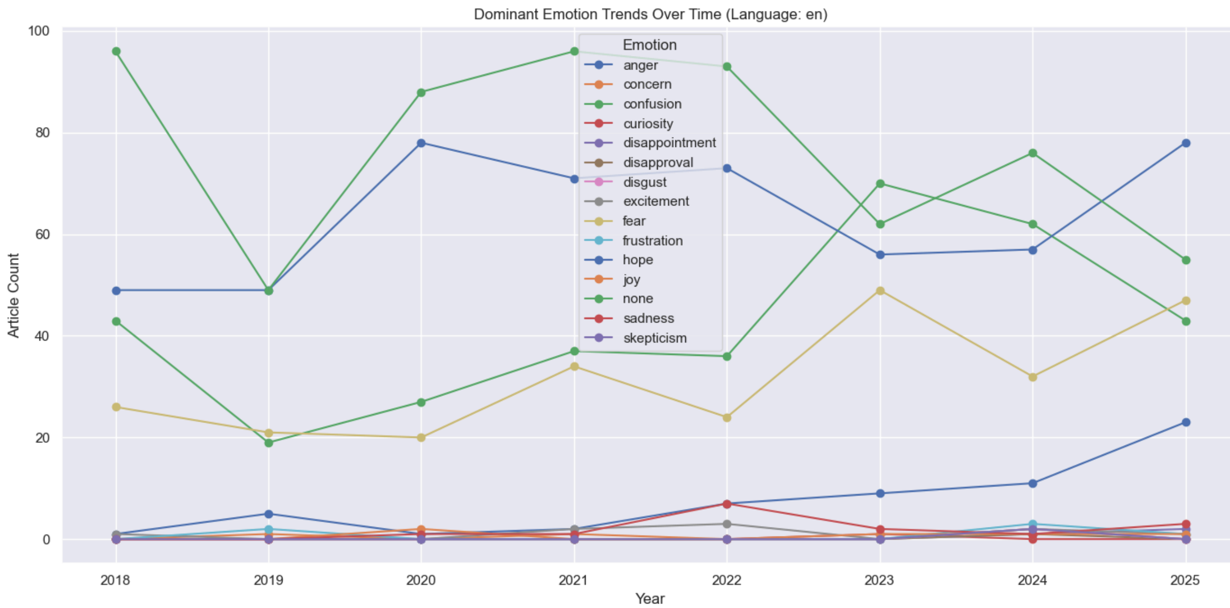


*Figure 14: Dominant Emotions Trends Over Time (English)*

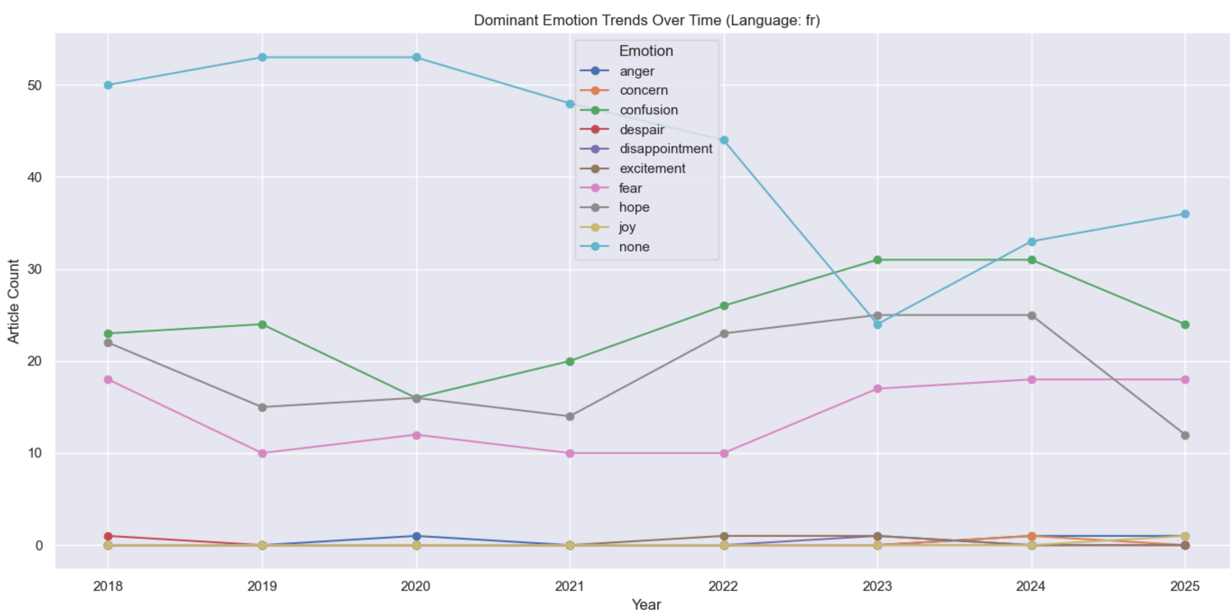


*Figure 15: Dominant Emotions Trends Over Time (French)*

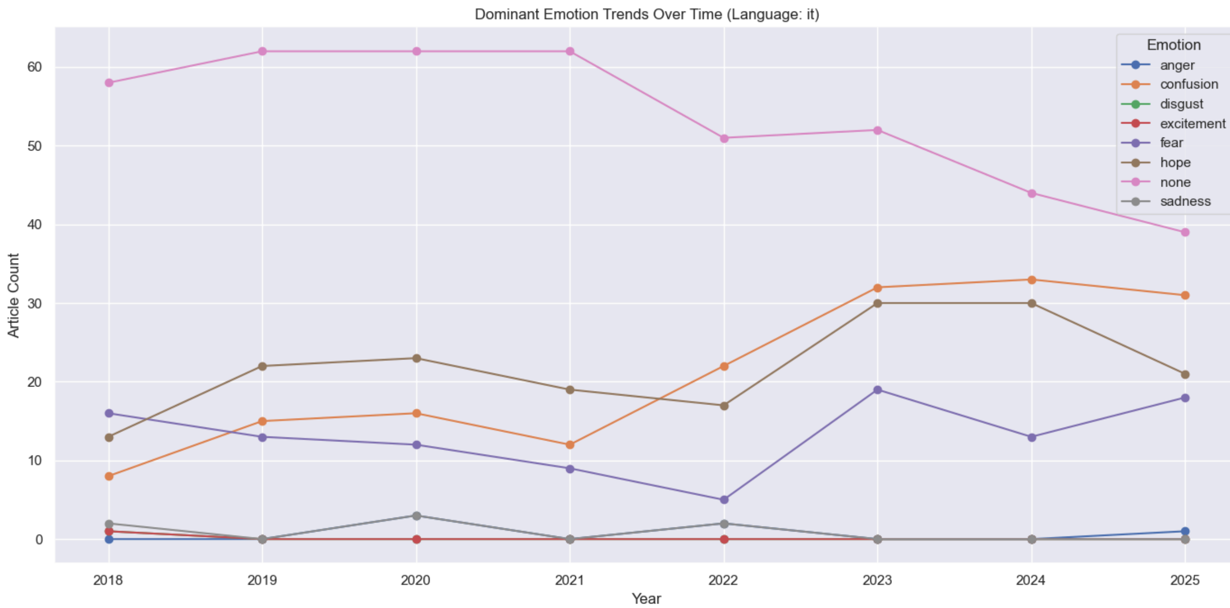


*Figure 16: Dominant Emotions Trends Over Time (Italian)*

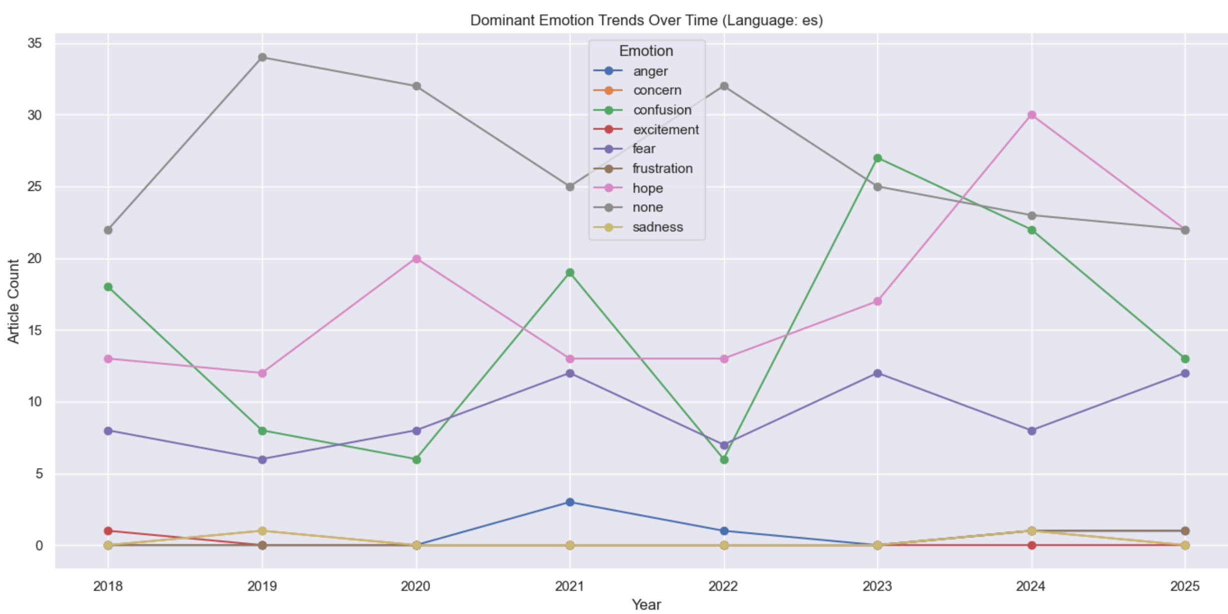


*Figure 17: Dominant Emotions Trends Over Time (Spanish)*