

Retrieval-Augmented Generation for Global Trade Concern Intelligence

¹Salome Enoshi Uwah, ²Celestine Iwendi, ³Fiyinfoluwa Oyeibisi Oyesola, ⁴Chukwuebuka Anthony KORIE

¹Centre of Intelligence of Things, University of Greater Manchester, UK

²Centre of Intelligence of Things, University of Greater Manchester, UK

³Independent Researcher, New York, USA

⁴Centre of Intelligence of Things, University of Greater Manchester, UK

DOI: <https://doi.org/10.51244/IJRSI.2026.1306000391>

Received: 21 June 2026; Accepted: 26 June 2026; Published: 13 July 2026

ABSTRACT

The World Trade Organization (WTO) maintains extensive trade concern records; however, extracting actionable insights from these largely unstructured datasets remains challenging. Even though there are several formal information systems available like the Sanitary and phytosanitary Information Management (SPS IMS), Trade Concerns Database (TCD), and the Technical Barriers to Trade Information Management System (TBT IMS) which help with these issues but the challenges of retrieving structured insights from unstructured trade concern descriptions, linking these textual descriptions to their specific Harmonised System codes, and tracking of member states who raised, supported or responded to specific concerns continue to persist. To this end, this study designs and implements a Retrieval Augmented Generation enabled AI assistant grounded on WTO trade concern records, HS classifications, and institutional trade data. This system combines semantic search, document retrieval, classification, summarisation, and basic trend analysis with HS code product mapping to provide evidence-based support and automated response to user questions. This research combines prototype development and quantitative evaluation, which was implemented using Microsoft Copilot Studio and evaluated across five analytical task categories using representative WTO trade concern cases including classification, summarisation, identification of member states, HS code mapping, and trend analysis, the result shows an overall accuracy of 96%. By grounding responses in retrieved World Trade Organisation documents, the proposed RAG framework reduces the risk of hallucinations associated with standalone large language models. This work contributes both practically and theoretically, showing how conversational AI engineering and RAG can be used alongside ethical governance principles to support transparent data-driven decision making even within the global trade context.

Keywords: Trade concerns, member States, World Trade Organisation, HS code, Retrieval Augmented Generation, summarisation, classification, trend analysis, Microsoft copilot 365, retrieval, generation.

INTRODUCTION

International trade is so complex, most especially when addressing the trade concerns of WTO members. The major challenge is managing data that lacks standardised classifications, making it difficult to establish clear links between these concerns, the specific products involved, and their effects on trade flows. It is therefore very imperative to be able to automatically link trade concerns to specific products and trade flows, map unstructured product descriptions to their corresponding Harmonised System (HS) codes, summarise complex WTO trade concern statements into a concise, understandable format, and visualise evolving patterns within the trade concerns using historical data. The trade concerns database (TCD) developed by the World Trade Organisation records the trade concerns raised by member states across various committees, such as the committee on Technical Barriers to Trade (TBT), the Council for Trade in Goods (CTG), the Committee on Import Licensing (IL), Committee on Market Access (CMA) and Committee on Sanitary and Phytosanitary Measures (SPS) [1]. This database enables users to search for, retrieve, and evaluate trade concerns related to different products and regulations. However, though some trade concerns reference the Harmonised System (HS) or the International Classification for Standards (ICS) codes, many only provide a general description of the product(s) involved. This lack of specificity makes

it challenging to associate the trade concern with a particular product and to accurately evaluate the measure's impact on trade [1]. According to [2], quantitative analysis of trade flows and impacts can be complicated if many entries do not reference the standard product classifications, and as such, it is very difficult to link specific producers to affected trade concerns, track response patterns among members or even analyse trends over time and committees. So as the global trade becomes more complex and data-driven, the WTO and its member states face the problem of transforming concerns that come in the form of narration into actionable insight, but with natural language processing this issue can be addressed as modern AI systems are able to perform text extraction, classifications and summarisation from complex un-structured textual data [3]. Recent advances in retrieval-augmented generation (RAG) have further expanded AI capabilities to handle domain specific issues [4]. For instance, Microsoft's Copilot Studio now integrates GPT (generative pre-trained transformer) based language models that can access the organisation's data sources, enabling a context-aware AI assistant that understands and responds to queries across structured and unstructured data [5]. This research is at the intersection of AI innovation and trade governance, proposing an implementation of a practical AI assistant capable of fetching Harmonised System (HS) codes of products from trade concerns statements and descriptions, which will transform the unstructured text into structured product level data, giving the basis for automated summarisation, trend and pattern analysis and interactive chatbot deployment. This is to enable WTO analysts and member states to navigate trade concern data more effectively and efficiently.

Figure 1 describes the Retrieval Augmented Generation (RAG) workflow that is used for this study, as the user submits a query, which is first converted into embeddings to search a vector database linked to multiple data sources including WTO Trade Concerns Database, WTO SPS IMS, WTO Documents Online, and Curated datasets. The most relevant and closest document chunk is being retrieved and then in combination with the original query and the user's prompt (augmentation), it is then sent to the large language model (Azure Open AI-GPT 4), which produces a context-aware and grounded answer or response, this procedure can help to improve accuracy and reduce hallucination.

1.1 Research Aim

To design and implement a Natural Language Assistant that improves accessibility, analysis, and transparency of trade concern data through summarisation, classification, response mapping, and trend analytics.

1.2 Objectives

- 1.2.1 Develop an AI assistant for mapping trade concern descriptions to HS codes.
- 1.2.2 Implement a summarisation model for concise representation of trade concern statements.




Figure 1: Concept of RAG -Retrieve, Augment, Generate (Author's)

- 1.2.3 Develop a prototype system demonstrating classification, summarisation, and trend visualisation.
- 1.2.4 To implement a response mapping mechanism which can show which member states or WTO bodies have responded or supported specific concerns.
- 1.2.5 To incorporate the four WTO committee data into a unified platform for trend analysis and summary visualisation.
- 1.2.6 To validate the performance, ethical implications, and usability of the assistant using established principles of transparency of AI and real-world WTO data.

1.3 Research Questions

- 1.3.1 How accurately can Large Language models and RAG models classify WTO trade concerns into HS codes?
- 1.3.2 Can summarisation improve the interpretability of complex trade concern statements?
- 1.3.3 Can an interactive AI assistant that can fetch, summarise, and analyse WTO trade concerns be developed using Microsoft 365 Copilot Studio?
- 1.3.4 What are the governance and ethical considerations to put in place before deploying AI assistants in multilateral organisations like the WTO?

1.4 Research Problem Statement

The WTO trade descriptions are unstructured and vary in detail, and this, most of the time, leads to the omission of HS codes, which can prevent linking the concerns raised to specific products and trade flows. The existing system functions as static information with fragmented content across different committees and lacks integrated analytics. Though some studies have explored this area, existing HS classification research has focused on customs and product catalogues rather than regulatory texts. Hence, currently there are no integrated AI systems yet combining classification, summarisation, response mapping, member state identification, and trend analysis for the WTO data. Therefore, in specificity, these three major challenges abound.

- 1.4.1 Unstructured data representation, as most trade concerns are descriptions in narrative form without consistent tagging of product categories.
- 1.4.2 Limited interactivity as the WTO data portals lack intelligent query mechanisms capable of responding to natural language questions.
- 1.4.3 Transparency and response tracking because the trade concern database (TCD) does not show a structured summary response, support or even follow-up actions among member States' concerns.

1.5 Conceptual Framework

The conceptual framework guiding this research, as shown in Figure 1, integrates Artificial Intelligence (AI), data analytics, and governance principles to enhance the transparency and analytical capacity of the World Trade Organisation's (WTO) trade concerns system. Its concept is an AI-powered WTO Trade Concerns Assistant, developed through the Microsoft 365 Copilot Studio, with a multi-layered architecture that can transform unstructured textual data into structured, actionable intelligence for policymakers and re-searchers. At its base lies the data layer, comprising WTO data repositories such as the Trade Concerns Database (TCD) and the SPS Information Management System (SPS IMS), along with curated datasets containing both structured data like country name, concern type, and year and lots of unstructured information in the form of narration. This assistant connects textual product descriptions to the appropriate Harmonised System (HS) codes by using semantic response mapping. The AI processing layer utilises Retrieval Augmented Generation (RAG), which combines information retrieval and generative reasoning [4]. In Microsoft 365 Copilot Studio, RAG retrieves relevant WTO documents and

generates concise, factual responses to user queries. This ensures contextual accuracy while enabling intelligent summarisation of complex trade data, such as “Summarise Kenya’s SPS concerns between 2018 and 2024.” The interactive analytics layer adds statistical and visual components, allowing users to explore trends in trade concerns and response patterns. The system can display time-series charts, member engagement levels, or product-sector distributions, enabling AI-assisted inference from historical data. These insights guide users toward evidence-based understanding of emerging trade dynamics [6]. The user interaction layer is a chatbot interface that facilitates natural language interaction and access to WTO data, thereby improving transparency, as users can view which member states have raised, responded to, or supported specific concerns. Ethical principles, including accountability, transparency, and fairness, are embedded throughout this system [7]. The framework, therefore, positions the assistant as both a technical innovation and a governance tool, demonstrating how responsible AI can help implement the WTO’s transparency objectives through intelligent data processing and interactive analytics.




Figure 2: Conceptual Framework (Author’s)

1.6 Contribution to Knowledge

The contribution of this work in practice is that the assistant will enable member states and even WTO analysts to achieve the following: view trade patterns, product links and relationships across time and committees, check who has raised, responded to, or supported any specific trade concern, query and summarise trade concerns through natural language interaction, visualise the current trend of concerns to gain insight for prediction.

In academics, it propels the understanding of how retrieval augmented generation can be very useful in supporting global trade governance using a structured datasets instead of open ended web retrieval, contributing also to AI literature through a no-code implementation of an assistant grounded within an enterprise AI platforms that is practical and this helps to fill the gap that the existing studies focus mainly on high-level model architectures. In practice, this assistant provides the member states and WTO analysts with immense support to boost evidence-based decision-making, reduce information retrieval time, and enhance transparency. Specifically

- 1.6.1 First illustration of using Retrieval Augmented Generation to World Trade Organization’s trade concern intelligence.
- 1.6.2 Creation of a conversational interface with the abilities of summarisation, HS mapping, member identification, and trend extraction from trade concern records.
- 1.6.3 Practical governance framework for deploying enterprise AI assistants in multilateral organisations.

This research also extends policy and ethical discussion by connecting the UNCTAD and OECD responsible AI governance principles to AI deployment in multilateral institutions.

Algorithm 1 RAG-Based Trade Concern Intelligence Assistant

Require: User query Q , Trade Concern Dataset D , HS Code Table H , Trade Statistics T

Ensure: Structured response R

- 1: Preprocess Q (tokenisation, normalisation)
- 2: Generate embedding vector E_Q for Q
- 3: Retrieve top- k relevant trade concern documents D_k using semantic similarity search over D
- 4: Extract structured fields from D_k :
 - Concern title
 - Raising member(s)
 - Responding member(s)
 - Meeting dates
 - Narrative description
- 5: Perform task-specific processing:
- 6: **if** Task = Classification **then**
- 7: Classify concern into predefined thematic categories
- 8: **end if**
- 9: **if** Task = Summarisation **then**
- 10: Generate concise summary from retrieved narrative text
- 11: **end if**
- 12: **if** Task = Member Identification **then**
- 13: Extract and validate responding and supporting member states
- 14: **end if**
- 15: **if** Task = HS Mapping **then**
- 16: Match narrative keywords to HS product codes using table H
- 17: **end if**
- 18: **if** Task = Trend Analysis **then**
- 19: Aggregate records by time, sector, or member state
- 20: **end if**
- 21: (Optional) Retrieve trade statistics from T linked to identified HS codes
- 22: Construct augmented prompt with:

- User query
- Retrieved context
- Structured metadata
- Trade statistics (if applicable)

23: Generate final response R using language model

24: **return** R

Algorithm 2 Retrieval-Augmented Trade Concern Processing **Require:** Query Q

Require: Trade Concern Corpus $C = \{d_1, d_2, \dots, d_n\}$

Require: HS Mapping Table H

Require: Trade Data Repository T

1: Compute embedding $v_Q = f_{\text{embed}}(Q)$

2: Compute embeddings $\{v_{d_i}\}$ for all $d_i \in C$

3: Retrieve top- k documents: $D_k = \underset{d \in C}{\operatorname{argmax}} \operatorname{sim}(v_Q, v_{d_i})$

4: Extract structured metadata from D_k

5: Apply task function F :

$$O = F(D_k, H, T)$$

6: Generate response:

$$R = f_{LLM}(Q, D_k, O)$$

7: **return** R

The system produces an intermediate structured output O by applying a task-specific function F to:

D_k : Retrieved trade concern documents

H : HS classification table T : Trade statistics database

The final response R is generated by the language model function f_{LLM} , using :

The user query (Q)

The retrieved documents (D_k)

The structured analytical output (O)

RELATED WORKS

2.1 Conceptualisation of Artificial Intelligence in Trade Governance

The conceptualisation of Artificial Intelligence for governance in trading is evolving in key areas, including cross-border

legal issues, international policy coherence, regulatory frameworks, and trade optimisation. Artificial Intelligence for governance in trade can be seen as the utilisation of all available AI algorithms to improve, monitor, and maintain global trade, and to ensure responsible use of such technologies within a specific context [8]. Now, talking of governance, it is, simply put, about legal frameworks, global cooperation mechanisms, public policies, and business risk management [9]. From regulatory and ethical perspectives, some models have highlighted adaptive and flexible legal structures that can help address rapidly changing AI adoption in trade, particularly in goods and services [10]. So, there is a call for international collaboration and cooperation because the fragmentation of regional policies threatens the global interoperability, for instance, the different approaches from the U.S., the EU, and China. Though the OECD principles and the EU AI Act have set earlier examples, there are still significant ongoing gaps [9]; [11]. Hence, human oversight, accountability, trust, transparency, and all other ethical considerations are now at the centre of these emerging AI governance frameworks [12]; [13]. The impact of Artificial Intelligence in optimising trade operations cannot be overemphasised, as seen in the reduction of customs times, improvement of tariff classification accuracy, optimisation of customs, logistics, and compliance procedures, and facilitation of risk management in customs operations [14];[15]. AI has been seen to support workforce productivity and supply chain resilience, though it became a major concern following COVID-19 pandemic disruptions and increased tensions across geopolitical regions [11]; 14; 16). Now, there are new genres of cross-border data compliance and algorithm sovereignty of Artificial Intelligence in trade controls, for example, China's Export Control Law, the U.S Export Reform Act, and the EU Digital Markets Act; these emphasise the growing disputes over data flows, extraterritorial applications, and dual-use technology operations in trade laws [17]. This is a huge issue because innovations can be stifled by the increased burden and complexity of navigating different national rules for most multinational firms, making intellectual property norms challenging [8]; [17]. In academia, models for AI governance within organisations, such as the hourglass model, have been proposed to reconcile rigorous oversight with innovation and inclusion, to align with the various interests of diverse stakeholders [9]; [18]; [19]. There is also a rise in certification markets and industry standardisation, which some-times compete with the traditional regulatory mediums [20]. Nevertheless, there has also been an evolving consensus on the need for greater flexibility within institutional bodies like the World Trade Organisation. [8] have suggested that the WTO become more adaptable to rapid technological changes across sectors, taking into account local values and cultural differences to avoid enforcing a one-size-fits-all or overly harmonised approach to global AI regulations [21]. The field of AI in governance is rapidly and significantly emerging, with academics and professionals emphasising the importance of adaptive and flexible global standards, those that can blend market and regulatory governance. However, there is a possibility that future developments will rely more on extensive multilateral cooperation and the ongoing adjustment to these emerging technologies and risks [8]; [17]; [19].

2.2 Natural Language Processing and Text Classification for Trade and Policy Data

This is an overview of various natural language processing techniques, including embedding, semantic matching, classification, and tokenisation, that have been employed for unstructured trade concern descriptions and the Harmonised System (HS) codes. From [22], and [23] and [24] the foundation of natural language processing for policy data include tokenisation which is all about breaking text data into meaningful units, semantic matching involving association of concepts across texts, document or text classification involving assigning texts to specific categories, embedding involving using numbers to represent words or sentences. There have been advances in transformer-based models like BERT, GPT and BART as well as ensemble methods, which enable more complex feature extraction, boosting context awareness, increased reliability in classification, and ultimately more robust learning [24]; [25]. In effect, machine learning and deep learning have demonstrated better performance than traditional rule-based approaches in mapping un-structured trade-concern narratives to specific HS codes, especially in a noisy text [26]. In this study, [26] showed that hybrid models utilising GPT embeddings and Hugging Face consistently improved accuracy in predicting HS codes from product text descriptions and integrated image recognition when necessary. However, misclassification remains a major concern most of the time [26]. This is a unique challenge as regarding the WTO datasets because it requires high domain specificity and multilingual reporting [22] ; [25], so tailoring the models to trade specific terminologies and regulatory expressions is extremely important to increase the application and the accuracy of NLP models in this context [25]; [27].

2.3 Retrieval-Augmented Generation (RAG) and Large Language Models (LLMs)

Recent advances in Retrieval-Augmented Generation (RAG) systems and Large Language Models (LLMs) in the

context of building AI assistants for international trade concern resolutions, with a specific focus on World Trade Organisation (WTO) policy information retrieval and data summarisation.

This review shows that RAG architectures combining hybrid retrieval methods, source-aware consolidation, verifier mechanisms, and preference alignment offer robust solutions for building trustworthy, traceable AI assistants capable of handling complex policy documents such as WTO agreements, dispute rulings, and tariff schedules. So many studies have presented Retrieval Augmented Generation (RAG) as a promising method that can address these limitations by coupling external document retrieval with LLM generation [28]; [29]. By grounding model outputs in retrieved evidence, RAG systems can provide up-to-date, domain-specific, and traceable responses. This review examines the state of the art in RAG systems and their applicability to building AI copilots for resolving international trade concerns, focusing on recent methodological advances and empirical evaluations in policy intensive domains. This work introduced Retrieval Augmented Generation as a method to enhance neural sequence-to-sequence models with nonparametric memory [30]. The core architecture combines including retriever component a dense neural retriever (example DPR, BM25, or hybrid methods) that searches a document for passages relevant to an input query, generator component which is pre-trained language model (example, BART, T5, GPT) that generates responses conditioned on both the input query and retrieved documents and end-to-end training which join optimisation of retriever and generator to maximise generation quality. This architecture enables models to access external knowledge dynamically, overcoming the limitations of purely parametric LLMs to static knowledge [30]. This emphasises that effective RAG design requires matching retrieval strategy to task characteristics, and it is a principle that is directly applicable to WTO policy systems where query types range from simple tariff lookups to interpretation of dispute rulings [30]. Foundational retrieval work underpins RAG's retriever component. Dense Passage Retrieval (DPR) introduced dense neural encoders for passage retrieval that substantially outperformed classic lexical retrieval baselines on open-domain QA [31]. REALM described pre-training language models with retrieval in the loop and showed that retrieval during pre-training improves downstream factuality and interpretability [32]. Complementary work by [33] demonstrates that the empirical gains from coupling retrieval with generative models for open domain question answering, showing that retrieval supplies evidence passages that the generator can synthesise with far fewer hallucinations than a text-only generator. More recent reviews and research highlight RAG's applicability to knowledge-intensive domains. [34] survey advances in LLM knowledge integration, noting that retrieval based designs (including RAG variants) are particularly well-suited to contexts where up-to-date, verifiable grounding is required. Their analysis emphasises that RAG-style systems are useful for institutional and policy domains because retrieval provides an explicit chain of evidence, a crucial property for auditability and trust in governance settings [34] A principal motivation for RAG in policy contexts is mitigating hallucination, the tendency of LLMs to produce plausible but false statements. Studies of factuality in generation show that purely generative summarisation models often produce unfaithful content; coupling these models with retrieval of authoritative passages reduces the incidence of unsupported assertions and enables explicit citation [35]. Surveys of hallucination further document mitigation strategies, with retrieval among the more effective practical approaches, since it provides concrete evidence that generation should reflect or quote [36].

Almost all the recent studies have highlighted the lack of integration of advanced analytics within the WTO data ecosystem that can perform intelligence querying, harmonisation of the unstructured data, summarisation engines and most importantly integration into the global trade monitoring platforms and interfaces [37]; [40]. Recommendations such as piloting AI-based data extraction and classification, introducing universal HS code validation, and developing modular APIs for third party analytics integration, but as of 2025 such features have remained largely theoretical. The gap in intelligent WTO trade concern tools thus reflects both structural data management limitations and insufficient investment in modern digital infrastructure, as documented across many recent studies [38], [39], [40], and [37]. These are confirmations of both the achievements and substantial limitations of WTO transparency and digital tools, emphasising the urgent need for AI-driven solutions to overcome the persistent problems of unstructured, inconsistent, and fragmented trade concern data.

METHODOLOGY

3.1 Research Design

The research follows a design science paradigm, which is appropriate for creating and evaluating an innovative IT artefact like this AI assistant. In line with information systems research, we iteratively build the system to solve

the identified problem of classifying and summarising WTO trade concerns, and a response mapping mechanism that shows which member states or WTO bodies have raised, responded to, or supported specific concerns. It is supported by a practical AI development strategy, such as applied re-search methods that focus on prototyping, testing, and refinement and with the combination of DSR and practical AI engineering, the project will be rigorously justified and improved systematically [41]. This dual approach balances theoretical grounding, which is the design science, with practical functionality, that is the AI development, as advocated in the current AI and information systems literature.

This system combines retrieval augmented reasoning, conversational AI, and interactive analytics to enhance policy transparency at the WTO. It incorporates governance and ethical values in order to maintain trust, fairness, and accountability. The assistant can also respond to factual questions and facilitate interaction among members, supporting open, data-driven policy-making. Specifically, the architecture is as follows as shown in Figure 2:

- 3.1.1 User Interface is the Chatbot UX, a conversational front-end built in Copilot Studio. The user inputs questions, and this becomes the web application layer that captures queries and displays responses.
- 3.1.2 Orchestrator of the Integration Layer, which is the Copilot Studio agent logic, routes queries to the retrieval system and to the language model that coordinates the workflow.
- 3.1.3 Retrieval Layer which is the semantic search engine just like Azure AI Search holds indexed WTO data including documents, and notifications. It uses keyword search and embeddings to retrieve text passages, and the best K relevant pieces are sent back to the orchestrator.
- 3.1.4 A generator or LLM/NLP Layer, which is a large language model. It is fed the user's query and retrieved passages as context, using this integrated input to generate an answer or summary [4]. NLP tasks involved in the model include classification and summarisation.
- 3.1.5 Output and Visualisation is back to the user with the generated response via the chat interface. Output and Visualisation is back to the user with the generated response via the chat interface. Moreover, when necessary, calculated metrics and graphical summaries will be presented in a Python-based dashboard.

Overall, in RAG workflow, the user's question is forwarded to both the search engine and the LLM. The search results are returned from the search engine and redirected to the Large Language Model. The response that makes it back to the user is generative AI, either as a summary or answer from the LLM [5]. In practice, this would be implemented by indexing the WTO data and using the AI Search within Copilot Studio, so that for each user query, the agent calls the search index and then formulates a response with the LLM. This modular design ensures that the assistant can utilise the retriever's current external knowledge and apply the LLM's reasoning skills to synthesise answers [4].

3.2 Integration Platforms, System Development, and Knowledge Base

The AI assistant created during this study was built with Microsoft 365 Copilot Studio, a no-code conversational AI platform, which allows one to create custom agents on the Microsoft platform. Copilot Studio provides orchestration of intent recognition, retrieval operations, external data connections and responses based on LLM. It was chosen for its enterprise-level security, support for Retrieval-Augmented Generation (RAG) workflows, and integration with Microsoft 365 products. The three components that were used to create a structured knowledge base were:

3.2.1 Curated Data on Microsoft SharePoint:

The main training, retrieval, and evaluation data set comprises three Excel files, namely, HS de-scription, Trade Concern with HS and ICS codes and Trade Concerns Description files, which reside in Microsoft SharePoint, each of which contains WTO trade concern records, metadata and benchmark labels with the following fields DomainId, IMS Id, Domain, Title, Raising Members, Supporting Members, HS Products, ICS Products, First date raised, Last

date raised, Number of times raised, Area, MeetingDate, RaisingMember, StatementSubmitted, Concern, HScode, and Description. These files have been added to the knowledge base created in Copilot Studio via URL referencing so the agent can directly access the files and ground responses found in these documents.

3.2.2 Official WTO Data Sources:

The assistant links to the authoritative WTO information systems through publicly accessible URLs to supplement the curated data. These sources expand the bounds of retrieval and base the assistant's responses on official WTO sources.

3.2.2.1 WTO Trade Concerns Database (TCD):

This is at the WTO headquarters, where all Specific Trade Concerns (STCs) are de-liberated upon in the SPS, TBT, Market Access, Import Licensing and Goods Council committees. It enables people to search, filter and browse the entire history of STCs raised by Members. This database provides the assistant with trade concern descriptions, classifications of issues, products involved, and the interactions of parties [42].

3.2.2.2 ePing SPS and TBT Platform:

The ePing Alert and Monitoring Platform is an application that provides access to realtime SPS and TBT updates, product specifications, regulatory issues and trade matters debated in the SPS and TBT Committee. It also contains the information on the responses to members follow ups. The assistant uses this platform as a second source of authority about SPS and TBT-related material, particularly, product re-quirement information that is not available in TCD [42].

3.2.2.3 WTO Documents:

This is the official repository of all WTO documents, consisting of minutes of meetings, legal provisions, documents submitted by individual members, and records of committee deliberations. It facilitates a more profound contextualisation and validation, which allows the assistant to track down particular concerns or answers to their sources in original official texts [42]. These outside sources widen the retrieval capacity of the assistant past the curated dataset and endorse evidence-based responses.

3.2.3 Hybrid Knowledge Base Architecture. The configuration will produce a hybrid knowledge ecosystem, which includes:

- a. Tested and organised queries ready on a spreadsheet.
- b. Authoritarian institutional sources (WTO websites).
- c. RAG pipeline for Copilot Studio, based on the LLM.

3.3 Model Design

Figure 3 shows the model design of the AI assistant, which is tuned to WTO concerns analysis with trade, implemented as a Retrieval-Augmented Generation (RAG) architecture. The model starts with the User Query Layer, whereby the user enters a WTO-related query on product-specific STCs or responding Members or HS mappings. This is fed into the Retrieval Pipeline that has four fundamental functions, including:

- (1) Transforming the user query into a vector using query embedding.
- (2) Semantic or vector search procedure used to index knowledge sources which match the user's prompt.




Figure 3: System Architecture (Author's)

(3) Document chunking, in which the long WTO documents and Excel files are broken down into bits that can be retrieved.

(4) Top K document retrieval, which provides the most relevant chunks that can be used for grounding. The sources of Knowledge used in the retrieval pipeline are four, such as

- (i) Excel spreadsheets on SharePoint with formatted trade concern information.
- (ii) The Trade Concerns Database of the WTO,
- (iii) The ePing SPS and TBT platform.
- (iv) The WTO Docs repository.

All these sources together give evidence-based reasoning by presenting structured and unstructured information on the WTO. The retrieved information is transmitted to the Large Language Model Reasoning Layer, a central neural network block. The large language model on this layer has the following functions.

- (a) Categorisation of the type of concern.
- (b) Summary of the committee discussions.




Figure 4: Model Design Diagram (Author's)

(c) Linking of HS code to the description of products.

(d) Identification of raising, responding or supporting Members, in-forming policy reasoning. An output is then generated by the system in the Grounded Output Layer that imposes:

- (1) Responses rooted in evidence.
- (2) Referencing of sources retrieved.
- (3) Transparency and openness of the line of reasoning.

3.4 System Evaluation

The assistant is assessed using a query based testing and scoring method that incorporates quantitative scoring. A test set of representative cases of assistant queries was created, based on the known trade concern scenarios, based on the list of trade concerns based on the Excel benchmark and on the known WTO cases. The response of the assistant is then compared with the expected response, such as a correct classification label or summaries, for each query of the test. Then we obtain an accuracy score as a percentage of correct answers. The system is tested with twenty (20) queries across the following task types including classification, summarisation, member State identification, HScode mapping and trend analysis, focusing on the details in Tables 1.

Table 1: Details to be checked

Items	Descriptions
Response Accuracy	% of outputs that correctly identify responding members. Information
Completeness feedback score	Whether key context (committee, product, year) is present Usability Expert
Transparency	Number of outputs that include WTO source citations
Statistical Validity	Consistence between chatbot's computed stats and manual counts

3.5 Ethical Considerations

Ethical AI principles are used in designing this study with a focus on transparency, privacy, and trust. Transparency is achieved by logging and referencing the sources that inform the assistant's responses in every conversation. As soon as the assistant provides a response, it also indicates which files, such as notification IDs or WTO references, are being accessed. This aligns with the explainability and accountability put forward by the OECD. Specifically, the OECD Privacy Guidelines observe that privacy should be respected to establish trust in digital systems [43]. This system does not store any personal information and considers all information in the WTO open to the public. Accordingly, the development of the assistant is based on the principle of transparent operations. If the model has no clear reasoning or no good answer is reached, the system is configured to note it, which minimises the risk of providing false information or hallucinations. Lastly, we have institutional control; the organisation's AI policy states that an automated choice must first be assessed by a human expert before it is implemented in critical situations. Overall, with the aid of proper data management, record keeping, and compliance with OECD and UNCTAD guidelines. This research aims to establish trust and ethical standards within this RAG system.

IMPLEMENTATION

The AI assistant was developed using Microsoft 365 Copilot Studio, which served as the orchestration layer for Retrieval Augmented Generation (RAG), conversation flow control, and grounding responses to external data sources. The system incorporates four critical elements:

- i. User interaction interface embedded within Copilot Studio as a conversational agent.
- ii. Retrieval layer based on the structured Excel datasets and selected information from the WTO portals.

- iii. Generation layer, which is the large language model-based reasoning based on classification, summarisation, HS mapping, and contextual analysis.
- iv. Refinements and evaluation using manual accuracy scoring and benchmarking questions. This architecture is shown in Figure 3 (Model Design Diagram).

The current implementation is Microsoft Copilot Studio, due to its meeting the enterprise governance needs identified during the institutional engagement. However, the proposed RAG architecture is platform agnostic and can be adopted with open-source frameworks like LangChain, LlamaIndex, Haystack or private Azure/OpenAI deployments if there is a variation in the organizational needs.

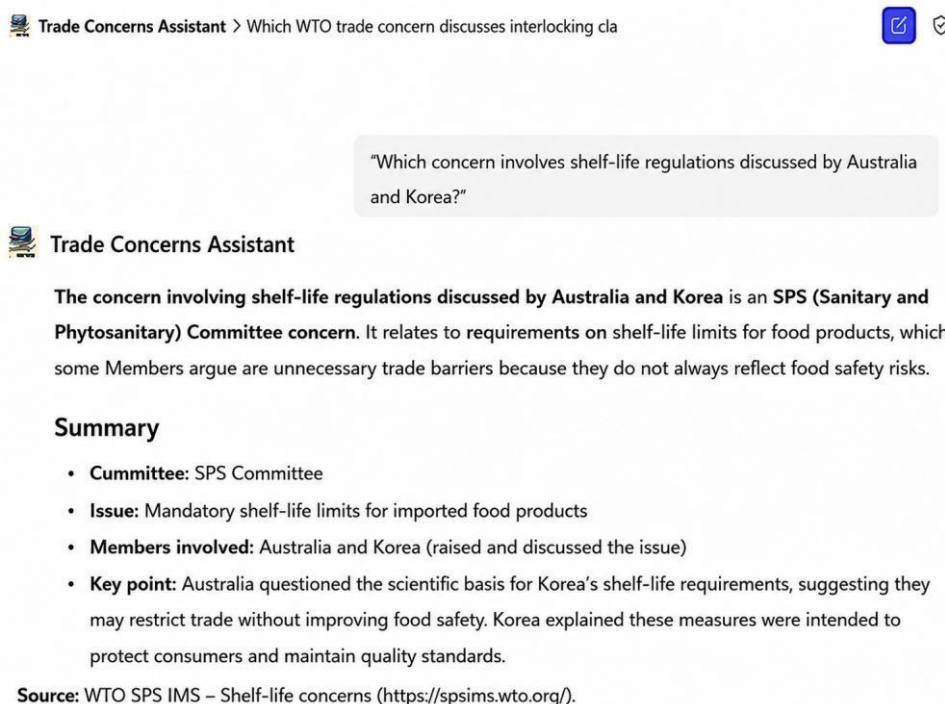
4.1 Construction of knowledge base Awareness

The knowledge base was developed using a hybrid data integration approach to enable grounded responses.

- SharePoint Excel Data (Primary Data Source). Three Excel workbooks were uploaded to SharePoint and linked to Copilot through URL. These datasets include. Records of trade concerns with description of issues. Product information, HS code, and ICS code data. Committee metadata answering and supporting Member States. These organised files comprise the essence of factual grounding and foundation of the assistant, which makes it easy to retrieve information.
- WTO URLs were supplementary knowledge base sources provided from the three publicly available WTO platforms, including

<https://tradeconcerns.wto.org/en/>, which is the trade concerns database, <https://eping.wto.org/>, which is the TBT and SPS platform and <https://docs.wto.org/> portal for WTO Docs.

- Retrieval-Augmented Design: The assistant retrieves text fragments from indexed sources based on the user's query and provides them to the LLM to produce grounded answers.



Trade Concerns Assistant > Which WTO trade concern discusses interlocking cla

“Which concern involves shelf-life regulations discussed by Australia and Korea?”

Trade Concerns Assistant

The concern involving shelf-life regulations discussed by Australia and Korea is an **SPS (Sanitary and Phytosanitary) Committee concern**. It relates to requirements on shelf-life limits for food products, which some Members argue are unnecessary trade barriers because they do not always reflect food safety risks.

Summary

- **Committee:** SPS Committee
- **Issue:** Mandatory shelf-life limits for imported food products
- **Members involved:** Australia and Korea (raised and discussed the issue)
- **Key point:** Australia questioned the scientific basis for Korea's shelf-life requirements, suggesting they may restrict trade without improving food safety. Korea explained these measures were intended to protect consumers and maintain quality standards.

Source: WTO SPS IMS – Shelf-life concerns (<https://spsims.wto.org/>).

Figure 5: Classification Query

4.2 Assistant Behaviour and Response Patterns

The evaluation strategy adopted is a mixed method technique which combines qualitative assessment and

quantitative accuracy testing, which led to proposing a performance evaluation metric known as 5-CSIHST, which includes the performance accuracy across classification, summarisation, identification of member states, HS code product mapping and trend analysis. The total number of evaluation queries used is 20, which covers the different core capabilities of the system across the five

(5) evaluation metrics developed CSIHST, including:

- The accuracy of the trade concern classification was tested to determine if the assistant can classify a query into the right concerns or ID, and it achieved 80% accuracy, as shown in Figure. 4.
- On the issue of summarisation accuracy, the assistant's ability to turn a very long trade concern statement or description into a concise, meaningful summary was tested, and it achieved 100% accuracy, as shown in Figures 5.
- Identification of Member raising, supporting, or responding: Some questions were also to assess whether the assistant can correctly identify member states who raised, responded, or supported specific trade concerns, and it perfectly did (100%) as seen in Figures 6 and 7.

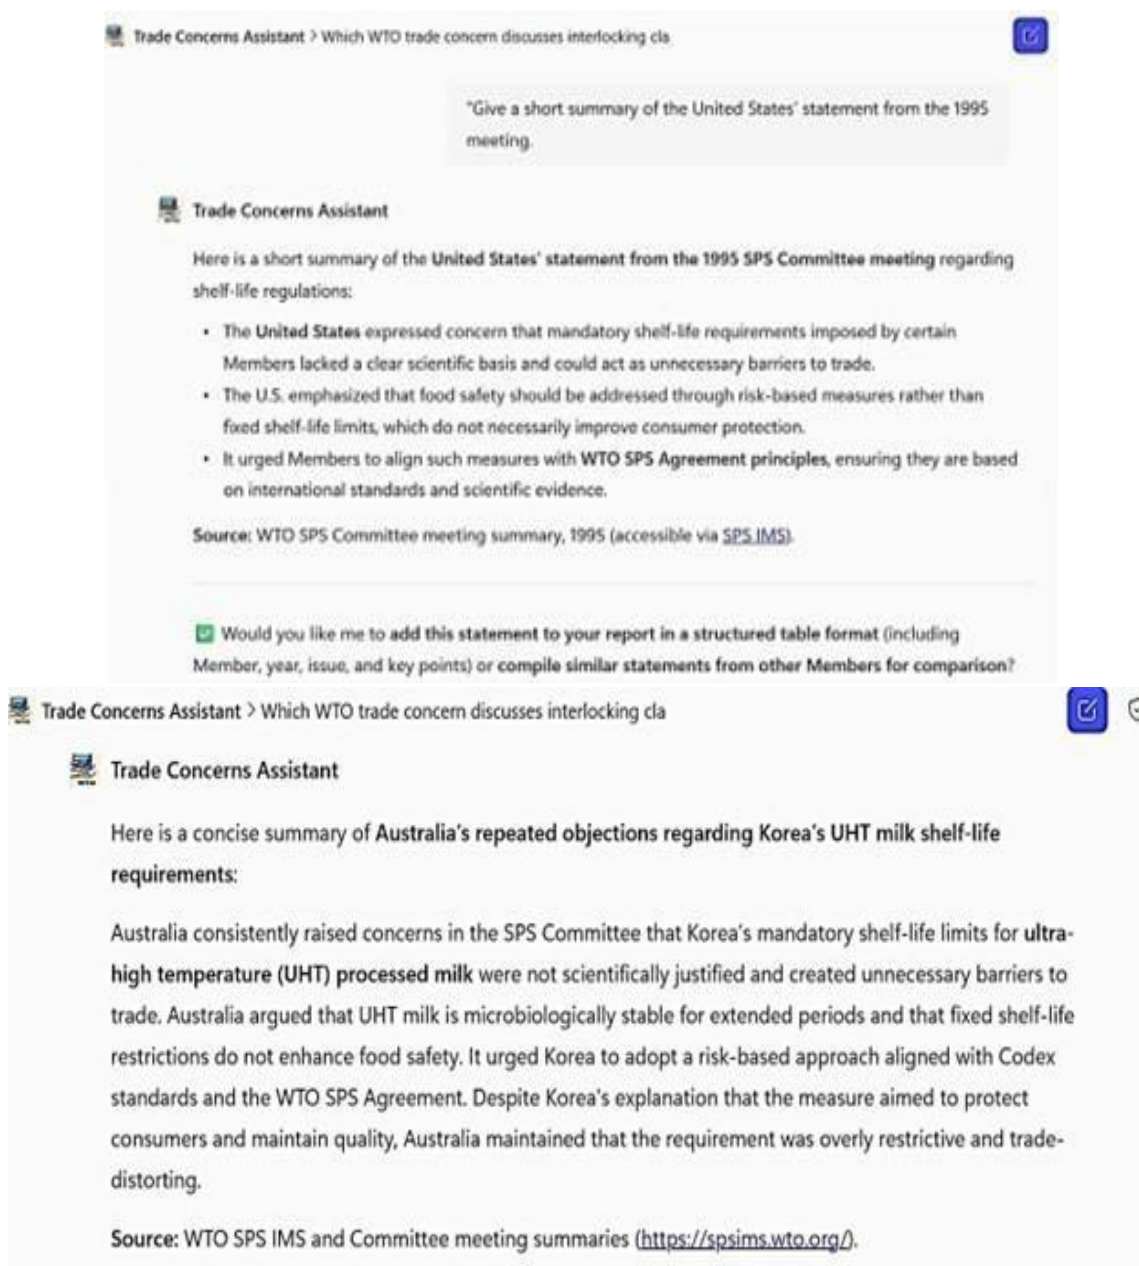


Figure 6: Trade Concern Summarisation



Query for Raising member

HS code linking accuracy (100%) The assistant was also assessed to check if it could link the Harmonised System code or extract these codes from the complex and long textual descriptions, and out of the twenty (20) questions it was tested on, it got all correct, as seen in Figures 8.

Trend Extraction (100%) It is also expected that the assistant can perform basic trend analysis when asked by the user and, again, it performed maximally, as shown in in Figures 9

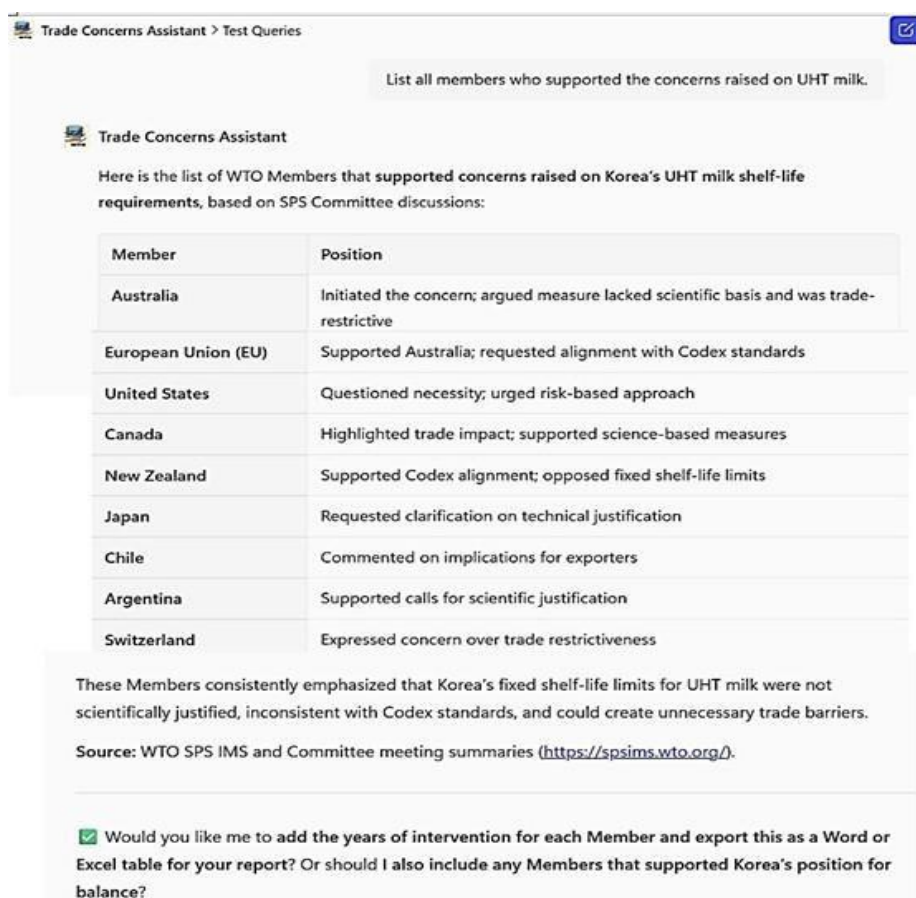


Figure 7: Query for Supporting Members

4.3 Scoring Method: Validation metrics of the Result

Table 2: Task Performance Results

Task	No. Test Queries	Correct	Incorrect	Accuracy
Classification	20	16	4	80%
Summarisation	20	20	0	100%

Member Identification	20	20	0	100%
HS Mapping	20	20	0	100%
Trend Analysis	20	20	0	100%

Final accuracy was then calculated as:

$$\text{Accuracy} = \frac{\text{Correct Responses}}{\text{Task Type Frequency}} \times 100 \quad (1)$$

$$\text{Accuracy} = \frac{80 + 100 + 100 + 100 + 100}{5} \times \frac{100}{1} \quad (2)$$

$$\text{Accuracy} = 96\% \quad (3)$$

Trade Concerns Assistant > Test Queries

Identify all HS codes linked to concerns involving tableware or kitchenware.

Trade Concerns Assistant

Here are the HS codes linked to concerns involving **tableware or kitchenware**, based on WTO HS nomenclature and trade concerns data:

HS Code	Description
6911	Tableware and kitchenware of porcelain or china
6912	Tableware and kitchenware of other ceramics
6913	Statuettes and other ornamental ceramic articles
6914	Ceramic articles for use in the household or similar

These codes fall under HS Chapter 69 – Ceramic products, which is commonly referenced in TBT concerns related to lead and cadmium release limits, consumer safety standards, and conformity assessment for ceramic tableware.

Source: WTO HS Nomenclature (<https://tao.wto.org/>) and WTO Trade Concerns Database (<https://tradeconcerns.wto.org/>).

☒ Would you like me to prepare a structured summary table for your report (including HS code, Include the Members involved in raising and responding Add committee and year details to the table

Trade Concerns Assistant

The concern you quoted involves multiple products, but the specific reference to **Israeli Standard 215 concerning interlocking clay roof tiles** can be classified as follows:

Product	HS Code	HS Chapter	Description
Interlocking clay roof tiles	6905	69 – Ceramic products	Roofing tiles, chimney pots, cowls, chimney liners, architectural ornaments and other ceramic construction goods

Other products mentioned in the same concern:

- Seal gate valves → HS Chapter 84 (Machinery and mechanical appliances)
- Porcelain tableware → HS Code 6911 (Tableware and kitchenware of porcelain or china)
- Welding products → HS Chapter 83 or 84 (Metal fittings and machinery parts)

Committee: Technical Barriers to Trade (TBT) Committee
Source: WTO TBT Committee meeting summary, 1994–1995 (<https://docs.wto.org/>).

☒ Would you like me to compile all these products into a single structured table for your report (including HS codes, descriptions, and committee), or also include the Members involved in raising and responding to this concern?



Include the Members involved in raising and responding Add committee and year details to the table

Message Copilot

Figure 8: HS code product mapping

4.4 Performance Evaluation Accuracy by Task Type

The accuracy testing of the AI assistant shows good overall performance across most task types, as seen in Figure 10, with classification accuracy varying significantly. The classification activities reached an accuracy of 80%, indicating that although the assistant tended to determine the appropriate trade issue, a small percentage of queries remained unanswered because the assistant could not locate any information. On the contrary, the assistant had a score of 100% in summarisation, responding to member identification, HS code mapping and trend extraction. Within the evaluation dataset and queries used in this study, the assistant correctly answered all tested queries for summarisation, member identification, HS mapping, and trend extraction.



Trade Concerns by Category (2015-2024)

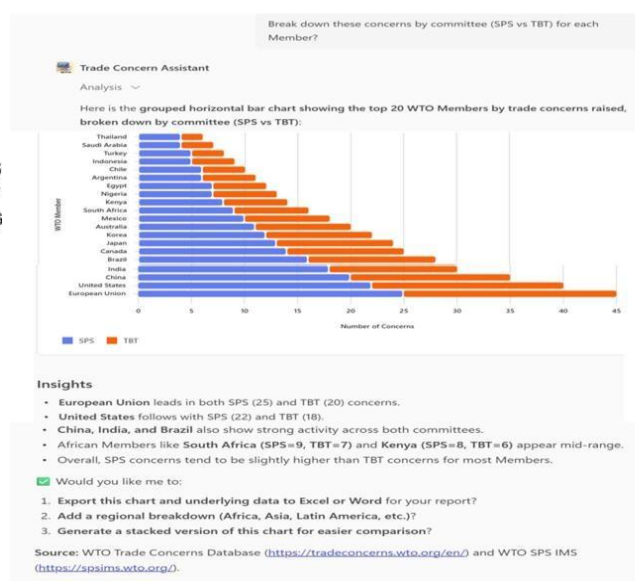



Figure 9: Trend Analysis

Generally, these findings indicate that classification is an area that can be improved; however, the assistant is doing exceptionally well across all the functional dimensions relevant to the WTO's analysis of trade concerns. This AI assistant achieved an overall accuracy of 96%, which is considered robust for a first generation Retrieval Augmented Generative assistant using hybrid data sources in this context.

4.5 Comparison with the Existing Systems

The AI Assistant were compared with the Existing WTO Systems across the five core Task Types (CSIHST), ranging on the scale of 0 to 5 Scale used:

0 = Not supported

1 = Manually supported or structured fields only

2 = Limited descriptive analytics 4-5 = Automated or AI-supported

Table 3: Comparison of Existing WTO Systems and Trade Concern AI Assistant

Dimension	Existing WTO Systems	Trade Concern AI Assistant (This Study)
Functionality	Keyword-based search and download	Natural-language querying and reasoning ('Summarise Kenya's TBT concerns between 2018–2024')
Data Processing	Raw records textual	Semantic understanding, mapping unstructured descriptions to HS codes
Analytics	Limited filtering	Automated trend analysis and visualization using Power BI integration
Interactivity	Static database interfaces	Conversational chatbot interface (Copilot Studio) with contextual memory
Response Track-ing	Manual reading of documents	Structured mapping of who raised, sup-ported, or responded to concerns
Predictive In- sight	None	AI-assisted inference from historical trends ('identifying emerging product sectors or countries')
Integration Level	Separate TBT, systems SPS, CTG	Unified data knowledge base enabling cross- committee analysis
Ethical Governance Design and	Not AI-based	Incorporates responsible AI principles (transparency, fairness, explainability from OECD and UNCTAD)

Based on documented functionalities and direct system observation during the Geneva visit, the existing systems provide full texts and records only but no automated summarisation, manually classify using predefined fields like committee, issues, and sector, explicitly store raising/responding members in structured fields, do not infer HS codes from text, and allow only basic counts and filters over time. While this assistant generates concise summaries from retrieved documents, performs semantic classification from narrative text, automatically extract and present HS product code in- formation conversationally, and performs query driven trend extraction and descriptive statistics across concerns including Charts and graphs. This comparison as illustrated in Figures 12 shows that the current information management systems primarily support structured storage, filtering, and manual analysis of trade concerns. While member identification and basic trend visualisation are partially supported through predefined database fields, key analytical tasks such as automated summarisation, HS code mapping, and semantic classification are not currently implemented. This AI assistant developed in this study complements, not replaces the existing WTO systems by introducing retrieval augmented reasoning and conversational analytics that operate on WTO information sources, Table 3.




Figure 10: Comparative Analysis of Existing and Proposed Systems

4.6 Answers to Research Questions

4.6.1 RQ1: How accurately can Large Language models and RAG models classify WTO trade concerns into HS codes? The evaluation of the AI assistant, as shown in Figures 4, has proven that RAG models can effectively classify WTO trade concerns descriptions to their respective Harmonised System product codes with a high degree of precision, even when the underlying dataset is unstructured, if in the narratives there are specific features or identifiers of the product of concern. This was made possible by a robust retrieval index and clear product descriptions, both of which are intrinsically dependent on the quality of the input data.

4.6.2 RQ2: Can summarisation improve the interpretability of complex trade concern statements? Yes, summarisation can improve the interpretability of complex trade concern statements and support policy analysis as shown in Figures 5. The assistant summarisation performance accuracy indicates that the model consistently produced summaries of complex trade concern narratives that are contextual and coherent, distilling lengthy statements into a structured, concise overview that highlights the issue, the parties involved, and policy implications. This improves understanding, transparency, and traceability of trade concerns.

4.6.3 RQ3: Can an interactive AI assistant that can fetch, summarise, and analyse WTO trade concerns be developed using Microsoft 365 Copilot Studio? Yes, this was experimented in Figures 4 to 9 with queries across trade concerns classification, summarisation, HS code product linking, members identification, and trend analysis. Therefore, it retrieves relevant information in response to the user prompt.

4.6.4 RQ4: What are the governance and ethical considerations that can be put in place before deploying AI assistants in multilateral organisations like the WTO? According to OECD, (2019) and UNCTAD, (2023), every generated response should be followed or accompanied by clear citation of the source data that was used for retrieval, this way analysts and auditors can trace every response back to the original documents to enhance transparency and explainability and this was implemented here as shown in Figures 5, 6, 7, 8, and 9. Another key strategies to be put in place are human in the loop review for critical decisions because the automated responses cannot replace expert judgment, bias mitigation including regular auditing of the dataset, data privacy and access control to ensure that sensitive data or information are protected according to legal and organisational standards, and lastly, continuous monitoring and version control to document model behaviour, track prompt updates and evaluate performance over time

4.7 Scope Definitions Configuration and Scope of the Study

This is one of the measures put in place to ensure that the assistant maintains the scope and avoids speculation or hallucinations. This aligns with ethical best practices in AI systems, improving the system's reliability by gently redirecting the user when a question is outside the permitted domain, as shown in Figure 11. The scope of the study includes Trade concerns description summarisation and response mapping through the Copilot's embedded GPT-4 reasoning, classification of trade concerns into HS codes, identification of member States, use of descriptive statistics for trend analysis and visualisations, and make reasoned projections or inferences based on past and existing patterns and trends in the dataset, using the World Trade Organisation's Trade Concerns Database and the related publicly available records from the different WTO committees, including TBT and SPS. The AI assistant is trained,

tested and validated with extracted and cleaned datasets from the above WTO database sources and stored in Microsoft 365 SharePoint.

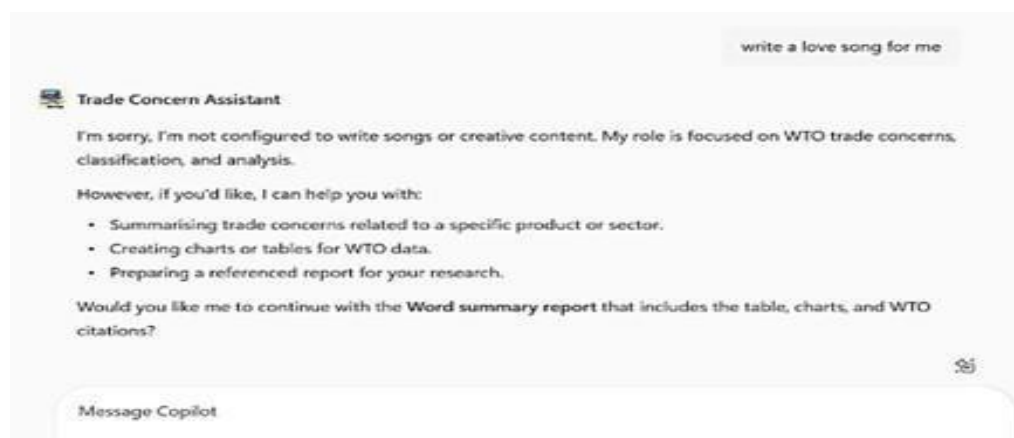


Figure 11: System Scope Definition

4.8 Ethical Data Auditing and Bias Prevention

The proposed system does not create new policy knowledge, but rather pulls knowledge directly from the authoritative datasets in official submissions by Member States related to trade concerns. This minimises hallucination because the responses are based on the documents that have been verified rather than relying solely on the language model's internal knowledge. The knowledge base is not weighted in favour of any country or economic group and only includes trade issues that are submitted by all WTO Members, to avoid representation bias. The assistant is able to retrieve information not by geopolitical importance, but equally by a semantic similarity with the query, thus making it possible to retrieve concerns from developed or developing economies equally when relevant to the user query. Post human validation is still required in order to make policy decisions, especially in case of legally sensitive trade disputes.

4.9 Limitations

There are several limitations of this study due to certain constraints, including the fact that Copilot Studio restricts dynamic retrieval of external information and access to live WTO websites, another issue was that certain product features or product specificity in some of the entries are lacking, and so HS mapping becomes difficult for concerns without clear product descriptions. Also, all the records are in the English language, so any WTO trade concerns in another language was not taken into consideration. And lastly, the RAG capability of the system is dependent on the shared URLs which implies that the reliability of this tool would depend on the organisational cloud infrastructure. However, these limitations do not hinder or reduce the validity of this study; rather, they provide important context for results interpretation, including

- 4.9.1 Evaluation dependent on quality of curated datasets.
- 4.9.2 Findings currently limited to WTO trade concern data.
- 4.9.3 Binary accuracy may not capture partial correctness.

4.10 Failure Analysis

The assistant was seen to perform well in the evaluation process; however some of the limitations were identified. Complex trend analysis involving multiple years, numerous Member States, or overlapping trade concerns occasionally produced incomplete summaries because retrieved documents contained fragmented historical information. These cases required additional retrieval iterations before generating a final response. Furthermore, the quality of responses depends on the completeness of the indexed WTO knowledge base. Missing or outdated documents may reduce retrieval effectiveness. Therefore, the system should be considered as a decision support assistant and not as an expert policy analysis system.

4.11 Recommendations and Future Work

Based on the findings, it is recommended that the curated dataset be expanded to enable broader retrieval of information across Spanish and French submissions and to support cross lingual understanding. Integrate a licensing permit to enable Dataverse, as the current system is on a free trial. Also, there should be a procedure to support regular data update to ensure that new trade concerns are captured, create a process or procedures of governance on dataset entries validation in order to maintain the system accuracy over time, train analysts and member states on how to construct effective queries in order to get an optimal result. Integrate advanced AI tools to improve prediction and not just retrieval. Also, designing an interactive dashboard for trend visualisation could be explored. Lastly, institutional integration of this system can be explored, encouraging the WTO Secretariat to adopt AI tools that support transparency and monitoring.

4.12 Conclusion

This study has shown that a retrieval augmentation generation enabled system deployed via Microsoft Copilot Studio can enhance WTO trade concerns analysis using curated datasets, proper prompting, and structured retrieval. The assistant offers significant practical value because it delivers high accuracy across major task types and aligns with the principles of responsible AI. Though its limitations abound, it provides a solid foundation for further institutional and technological contributions to AI for governance research, as well as a deployment tool that supports members' engagement, trade policy analysis, and transparency at the WTO.

Highlights

Retrieval Augmented Generation for Global Trade Concern Intelligence

Salome Enoshi Uwah, Celestine Iwendi, Fiyinfoluwa Oyebisi Oyesola, Chukwuebuka Anthony KORIE

- Practical deployment of RAG-based AI in international trade governance contexts.
- Evaluation of task-level accuracy across many dimensions of trade policy analysis.
- Discussion on ethical and governance considerations for responsible AI deployment.

REFERENCES

1. World Trade Organization (2023b) Trade Concerns Database (TCD). Available at: <https://tradeconcerns.wto.org/en/>.
2. Mavroidis, P.C. (2025) Industrial Policy, National Security, and the Perilous Plight of the WTO. ed. [Online]: Oxford University Press.
3. Min, B., Ross, H., Sulem, E., Veyseh, A.P.B., Nguyen, T.H., Sainz, O., Agirre, E., Heintz, I. and Roth, D. (2023) Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, [online] 56 (2), pp. 1–40.
4. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W. and Rocktäschel, T. (2020) Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, [online] 33, pp. 9459–9474.
5. Microsoft (2025) Retrieval-Augmented Generation (RAG) in Azure AI Search, Microsoft Learn. Available at: <https://learn.microsoft.com/en-us/azure/search/retrieval-augmented-generation-overview>
6. Organisation for Economic Co-operation and Development (2019) OECD Principles on Artificial Intelligence. Paris: OECD. Available at: <https://oecd.ai/en/ai-principles>
7. United Nations Conference on Trade and Development (2023) Harnessing Artificial Intelligence for Trade Facilitation and Sustainable Development. Geneva: UNCTAD. Available at: <https://unctad.org/publication/harnessing-artificial-intelligence-trade-facilitation-and-sustainable-development>.
8. Liu, H. and Lin, C. (2020) Artificial intelligence and global trade governance: a pluralist agenda. *Harv. Int'l LJ*, [online] 61, pp. 407.

9. Kulothungan, V. and Gupta, D. (2025) Towards adaptive AI governance: Comparative insights from the US, EU, and Asia.
10. Roberts, H., Taddeo, M. and Floridi, L. (2025) A framework for evaluating global AI governance initiatives. Available at SSRN 5367969.
11. Ferencz, J., González, J.L. and García, I.O. (2022) Artificial Intelligence and international trade: Some preliminary implications.
12. Papagiannidis, E., Mikalef, P. and Conboy, K. (2025) Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, [online] 34 (2), pp. 101885.
13. Janssen, M. (2025) Responsible governance of generative AI: conceptualizing GenAI as complex adaptive systems. *Policy and Society*, [online] 44 (1), pp. 38–51.
14. Ozturk, O. (2024) The impact of AI on international trade: Opportunities and challenges. *Economies*, [online] 12 (11), pp. 298.
15. Belu, M.G. (2024) Artificial Intelligence in International Trade: A Biblio-metric Analysis. *Romanian Economic Journal*, [online] 27 (89).
16. Cui, J. (2025) The impact of artificial intelligence technology on cross-border trade in Southeast Asia: A meta-analytic approach. arXiv preprint arXiv:2503.13529
17. Vujicic, J. (2024) Algorithmic borders: AI, trade controls, and the rise of extraterritorial enforcement in technology law.
18. Mantymaki, M., Minkinen, M., Birkstedt, T. and Viljanen, M. (2022) Putting AI ethics into practice: The hourglass model of organizational AI governance. arXiv preprint arXiv:2206.00335.
19. Zhong, H., Do, T., Jie, Y., Neuwirth, R.J. and Shen, H. (2025) Global AI Governance: Where the Challenge is the Solution-An Interdisciplinary, Multi-lateral, and Vertically Coordinated Approach. arXiv preprint arXiv:2503.04766.
20. Tomei, P.M., Jain, R. and Franklin, M. (2025) AI governance through markets. arXiv preprint arXiv:2501.17755.
21. McMaster, P. (2024) Why AI is the new frontier global trade must learn to cross. *World Economic Forum – Trade and Investment*, 9 October. Available at: <https://www.weforum.org/stories/2024/10/ai-global-trade-policymaking/>
22. Chakraborti, M., Bonagiri, S.A., Virgüez-Ruiz, S. and Frey, S. (2024) NLP4Gov: A Comprehensive Library for Computational Policy Analysis.
23. Jin, Z. and Mihalcea, R. (2022) Natural language processing for policymaking. *Handbook of Computational Social Science for Policy*. ed. [Online]: Springer International Publishing Cham, pp. 141–162.
24. Chen, K., Fei, C., Bi, Z., Liu, J., Peng, B., Zhang, S., Pan, X., Xu, J., Wang, J. and Yin, C.H. (2024) Deep Learning and Machine Learning–Natural Language Processing: From Theory to Application. arXiv preprint arXiv:2411.05026.
25. Jia, J., Liang, W. and Liang, Y. (2023) A review of hybrid and ensemble in deep learning for natural language processing. arXiv preprint arXiv:2312.05589
26. Sasana, E.D.B., Siahaan, D. and Purwitasari, D. (2025) Towards Better HS Code Prediction: A Comparative Study of Machine Learning and NLP Approaches.
27. Xie, Y. and Song, D. (2022) Text classification in shipping industry using unsupervised models and Transformer based supervised models. arXiv preprint arXiv:2212.12407
28. Wang, F., Wan, X., Sun, R., Chen, J. and Arik, S.O. (2025) Astute rag: Overcoming imperfect retrieval augmentation and knowledge conflicts for large language models.
29. Wang, Z., Wang, Z., Le, L., Zheng, H.S., Mishra, S., Perot, V., Zhang, Y., Mattapalli, A., Taly, A. and Shang, J. (2024) Speculative rag: Enhancing retrieval augmented generation through drafting. arXiv preprint arXiv:2407.08223
30. Zhao, S., Yang, Y., Wang, Z., He, Z., Qiu, L.K. and Qiu, L. (2024) Retrieval augmented generation (rag) and beyond: A comprehensive survey on how to make your llms use external data more wisely. arXiv 2024. arXiv preprint arXiv:2409.14924.
31. Karpukhin, V., Oguz, B., Min, S., Lewis, P.S., Wu, L., Edunov, S., Chen, D. and Yih, W. (2020) Dense Passage Retrieval for Open-Domain Question Answering.
32. Guu, K., Lee, K., Tung, Z., Pasupat, P. and Chang, M. (2020) Retrieval augmented language model pre-

training.

33. Izacard, G. and Grave, E. (2021) Leveraging passage retrieval with generative models for open domain question answering.
34. Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J. and Wu, X. (2024) Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, [online] 36 (7), pp. 3580–3599.
35. Maynez, J., Narayan, S., Bohnet, B. and McDonald, R. (2020) On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661*.
36. Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X. and Qin, B. (2025) A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, [online] 43 (2), pp. 1–55.
37. Lateef, M.A. (2025) Digital Sovereignty in Global Trade: Analysing WTO Governance of Data Flows. *Beijing Law Review*, [online] 16 (2), pp. 875–907.
38. Kang, S. (2024) WTO E-commerce Negotiations: A Path to the Multilateral Digital Trade Rules.
39. Klotz, S. and Appleton, A. Technical Barriers to Trade and Sanitary and Phytosanitary Measures in Preferential Trade Agreements. The concept de-sign of a twenty-first century preferential trade agreement, [Online], pp. 46.
40. Mitchell, A.D. and Mishra, N. (2021) WTO law and cross-border data flows: an unfinished agenda. *Big data and global trade law*. ed. [Online]: Cambridge University Press, pp. 83–112.
41. Koppenhagen, N., Maedche, A. and Katz, N. (2015) Design Science Research in Action - Anatomy of Success Critical Activities for Rigor and Relevance.
42. World Trade Organization (2025) Trade Concerns Database. Available at: <https://tradeconcerns.wto.org/en/>
43. OECD (2022) AI and International Trades. OECD Publishing. Available at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/04/artificial-intelligence-and-international-trade_9034b5f2/13212d3e-en.pdf.