



# AI-Powered SDMX Data Navigation with LLMs & JSON Schema

Alessandro Benedetti (Sease), Bianca Ligani (BIS), Rafael Schmidt (BIS)

SDMX Global Conference 2025 - 29/09/2025

# BIS Data Portal

## Overview

- BIS Data Portal launched in 2023
- Statistical Data and Metadata eXchange (**SDMX**) backbone
- Find out more typing in your browser [data.bis.org](https://data.bis.org)

## Key features

- ~ **1.2 million** time series published
- Multiple statistical domains (27 dataflows)

## Search capabilities

- **SOLR** search implemented within the portal
- **Synonyms** applied for common terms / abbreviations, currency-country, etc
- Weighted search across a number of facets:  
CATEGORY(^10), SERIES TITLE(^5), DIMENSIONS(^1), REFERENCE AREAS(^1)

## Overview

- 1 Use case review
- 2 LLM-based search for SDMX data
- 3 Structured output and JSON schema
- 4 Test framework
- 5 Future works

# AI dimensions filtering assistant

**QUERY:** What are the **office space prices** in **Rabat**

**Commercial Property Prices** ← **Topics Selected**

Overview | Tables & dashboards | **Data**

What are the office space prices in Rabat

Bookmark Share Export

**Country Selected**

**Dimension Selected: Covered Area**

**Dimension Selected: Real Estate Type**

**Filters**

Refine your data selection by specifying the criteria below.

Timespan:  Frequency: Quarterly  Reference area: Morocco  Covered area: Capital city/biggest city/...

Real estate type: Office & Retail  Real estate vintage: All  Compiling agency: Central bank  Priced unit: Pure price

Seasonal adjustment: Non seasonally adjust...

**Time series list**

1 found

Q.MA.2.D.0.0.6.0  
Morocco - Commercial property price index, office, retail buildings, Rabat

Index, 2006 Q1 = 100  
97,3

2005-Q1 - 2024-Q3  
Quarterly (Q)

Relevance

## BIS Data Portal AI assisted search - goals

Leverage **SDMX REST API** + **LLM** to assist in search process

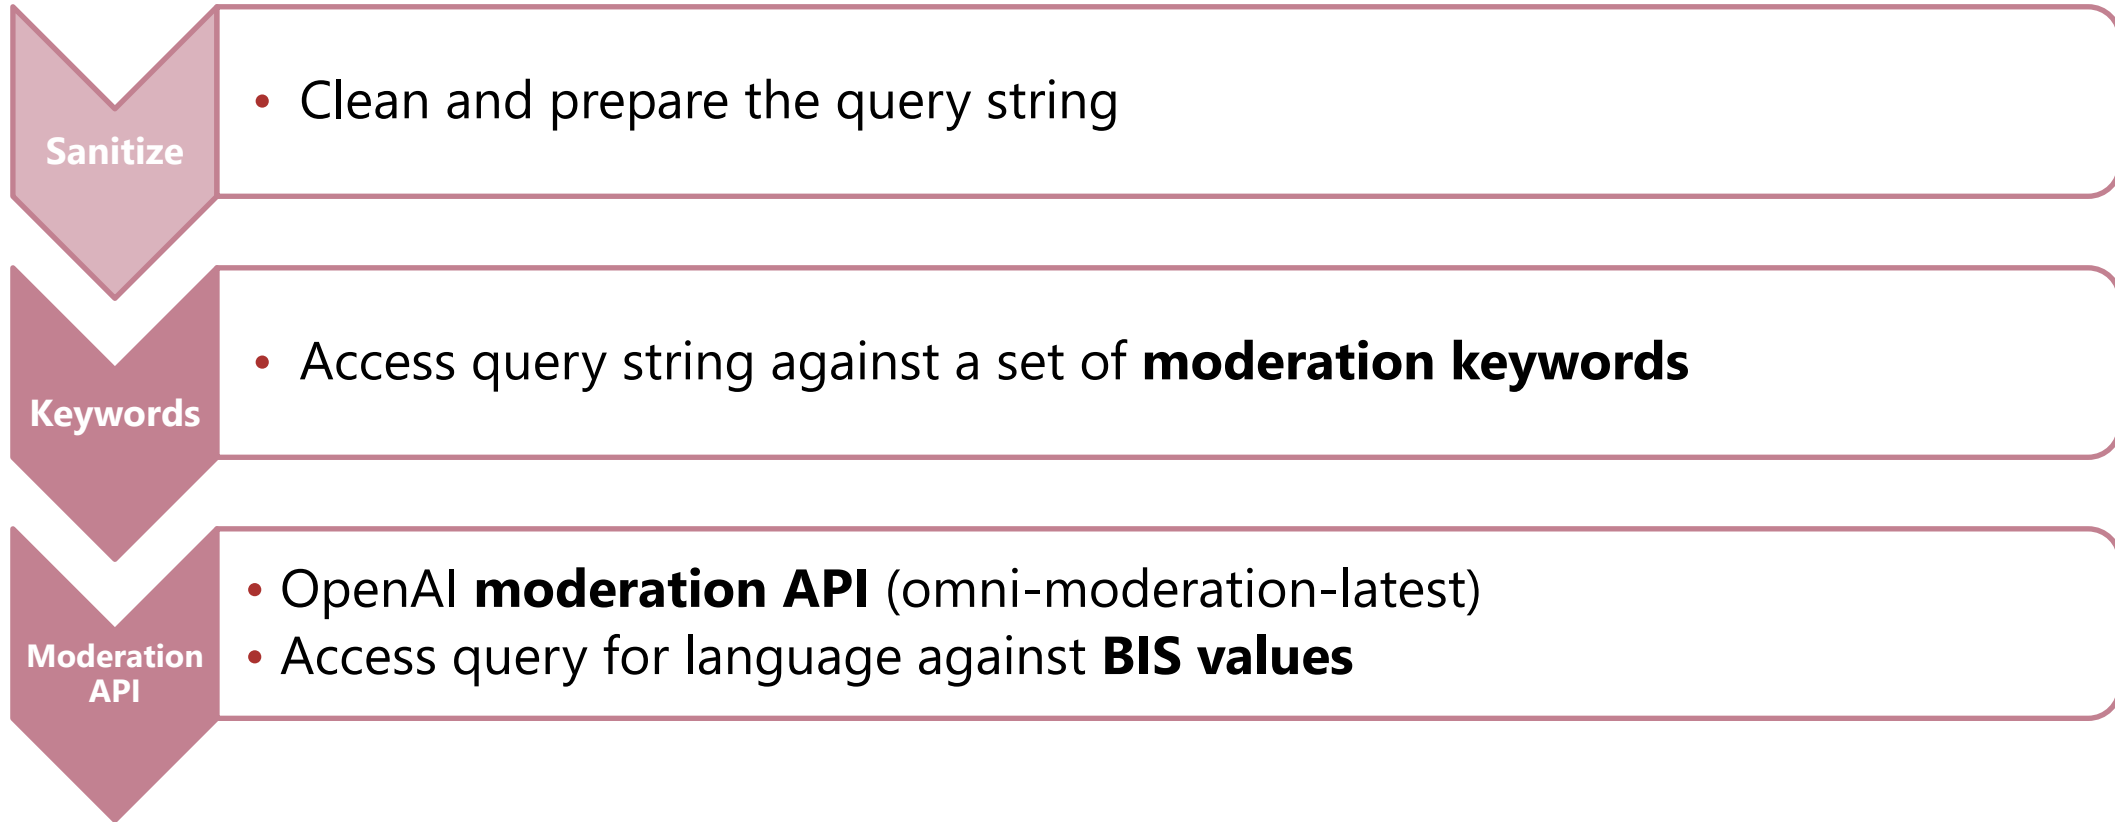
Allow for a *natural language* search within the BIS Data Portal

Provide search results where lexical search struggles

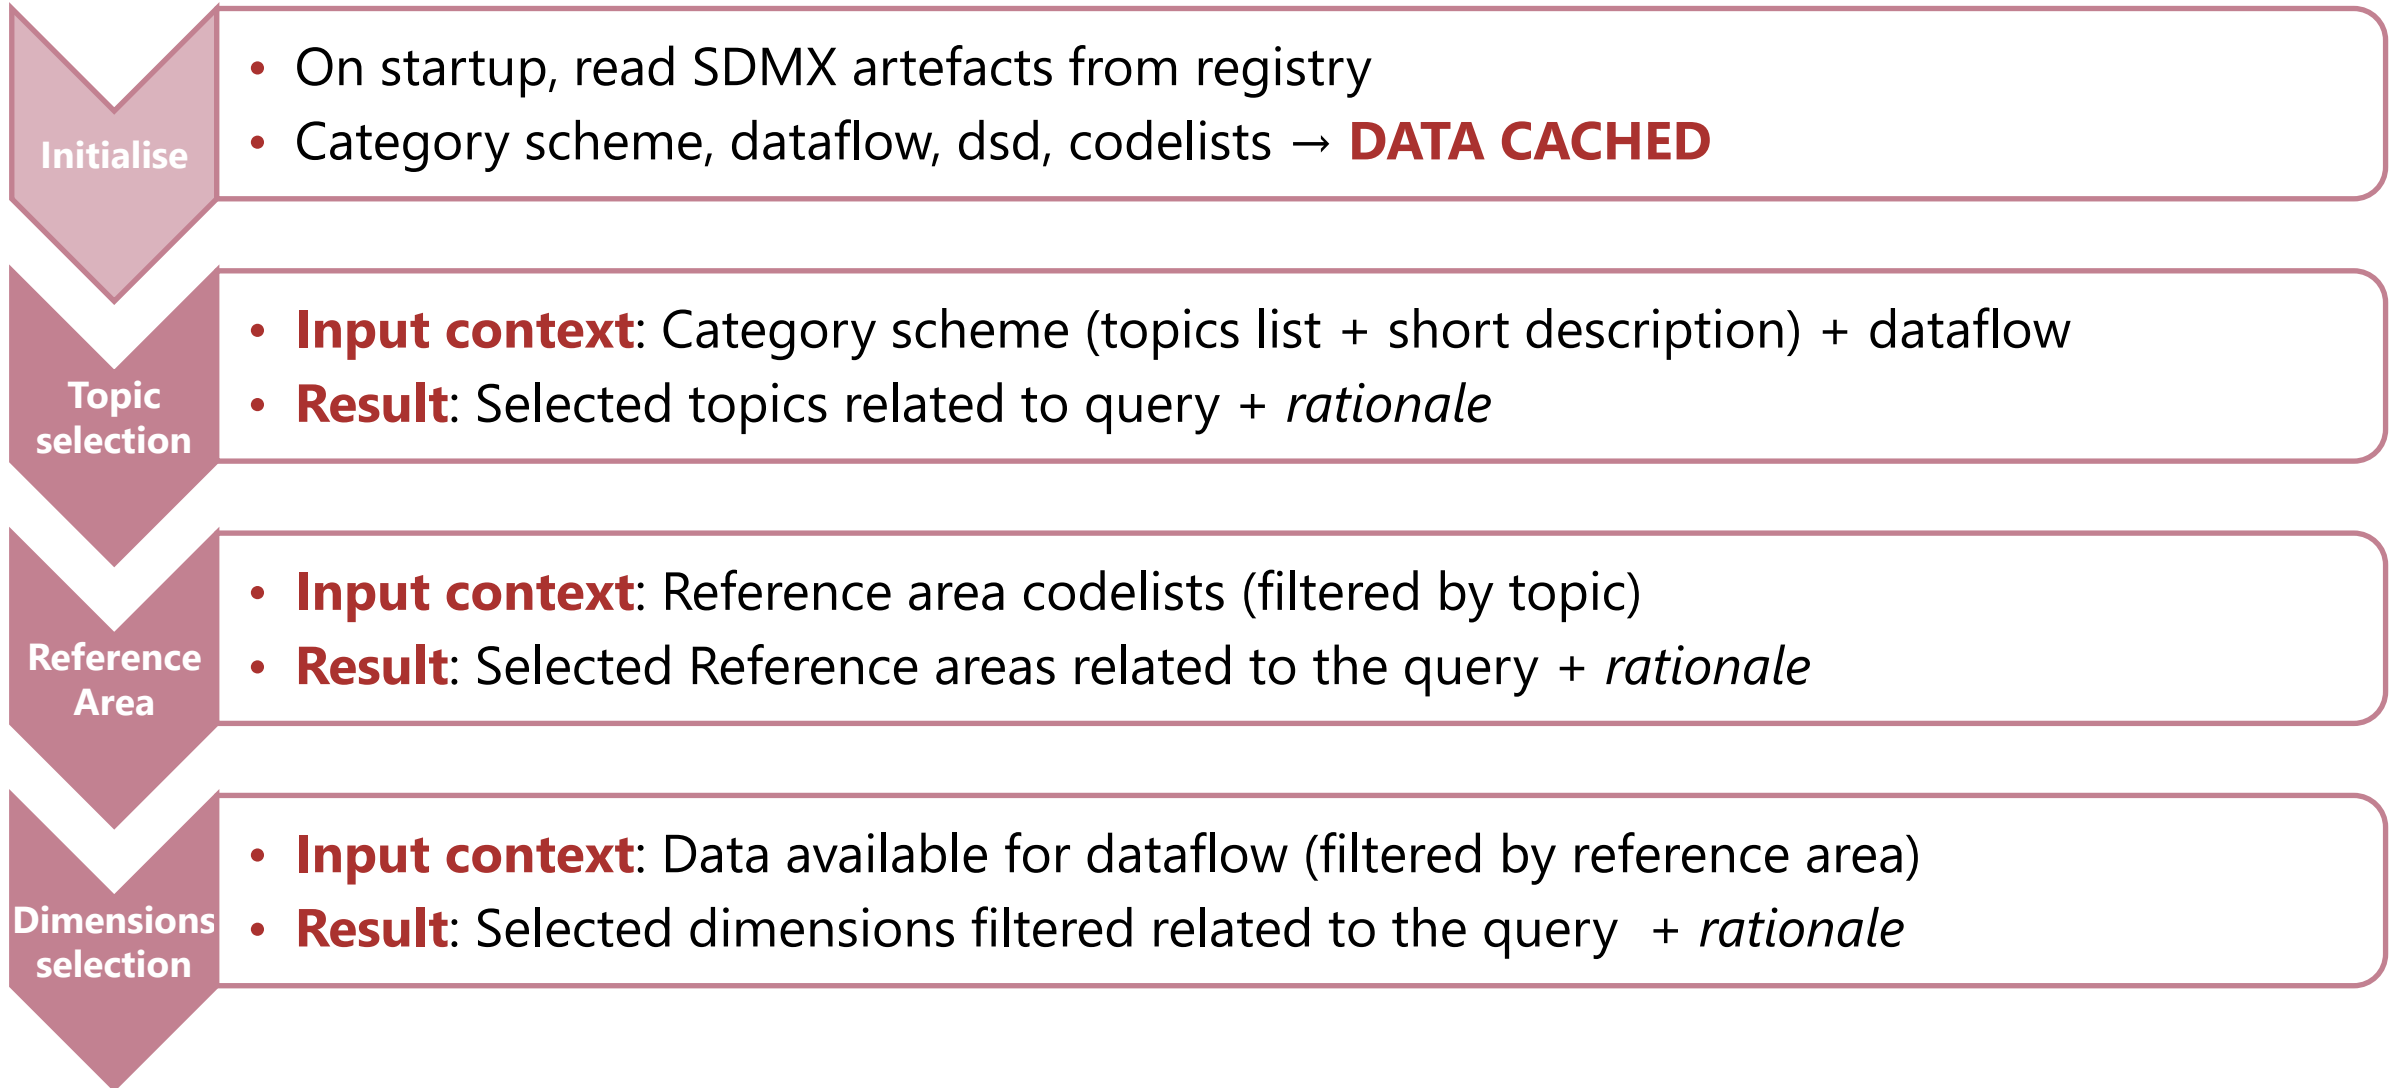
Provide rationale to end-user on why selections are being made

Provide a solution that can be shared with central banks / International Organizations within their data domains

# Query moderation



## Structured process flow





# Structured output - JSON schema

- **Structured output**: Enables the model to return responses in a **well-defined JSON format**.
- **JSON schema definition**: Specifies the **expected structure, data types, and constraints** of the output.
- **Controlled output**: The model generates responses that align with the predefined schema.

<https://openai.com/index/introducing-structured-outputs-in-the-api/>

```
"type": "json_schema",
"json_schema": {
  "name": "identified_geographical_areas",
  "schema": {
    "type": "object",
    "properties": {
      "geographical_areas": {
        "type": "array",
        "items": {
          "type": [
            "string",
            "null"
          ],
          "description": "List of geographical areas identified in the Query",
          "enum": ["Indonesia (ID)", "Türkiye (TR)", "Spain (ES)", etc..., "Euro area (XM)", null]
        }
      },
      "rationale": {
        "type": "string",
        "description": "Describe why the model identified these specific geographical areas in the Query"
      }
    },
    "required": ["geographical_areas", "rationale"],
    "additionalProperties": false
  },
  "strict": true
}
```



## Future works

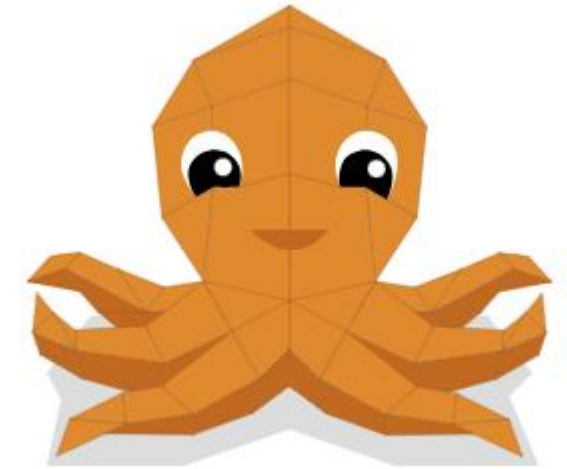
- **Few-shots** integration
- Potentially undertake **our own fine-tuning/distillation** for the specific task
- **Enriching the test framework query sets** by exploiting both LLM and domain experts for a more robust evaluation
- **Collect feedback** from both internal and public users
- **Refine prompts** for BIS published statistics
- **Open-source** the product for other CBs, NSIs and IOs to use as reference



Demo

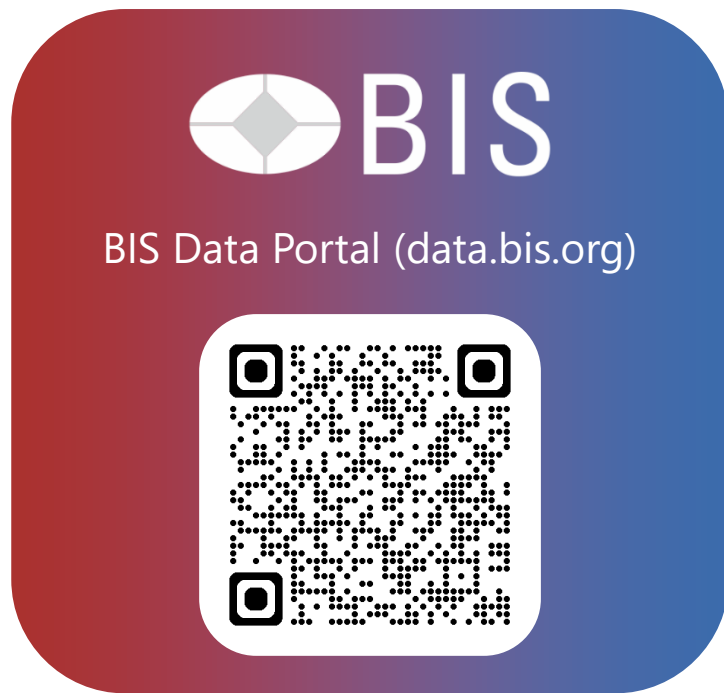
## sdmx.io – Solutions for official statistics use cases

- The *AI-assisted search* can be **integrated with any SDMX compliant webservice**
- The project team will prepare the **code as an open-source product** for CBs, NSIs and IOs to use as reference
- Serves as a **reference implementation** for organizations to **adapt and integrate the AI assisted search functionality into their own systems**
- **Promote collaboration** on software solutions in the SDMX community



*"The SDMX and official statistics community produces a lot of great open source software tools. sdmx.io brings these together to help solve concrete problems, with practical guidance and worked solutions to download and deploy."*

# Thank you!



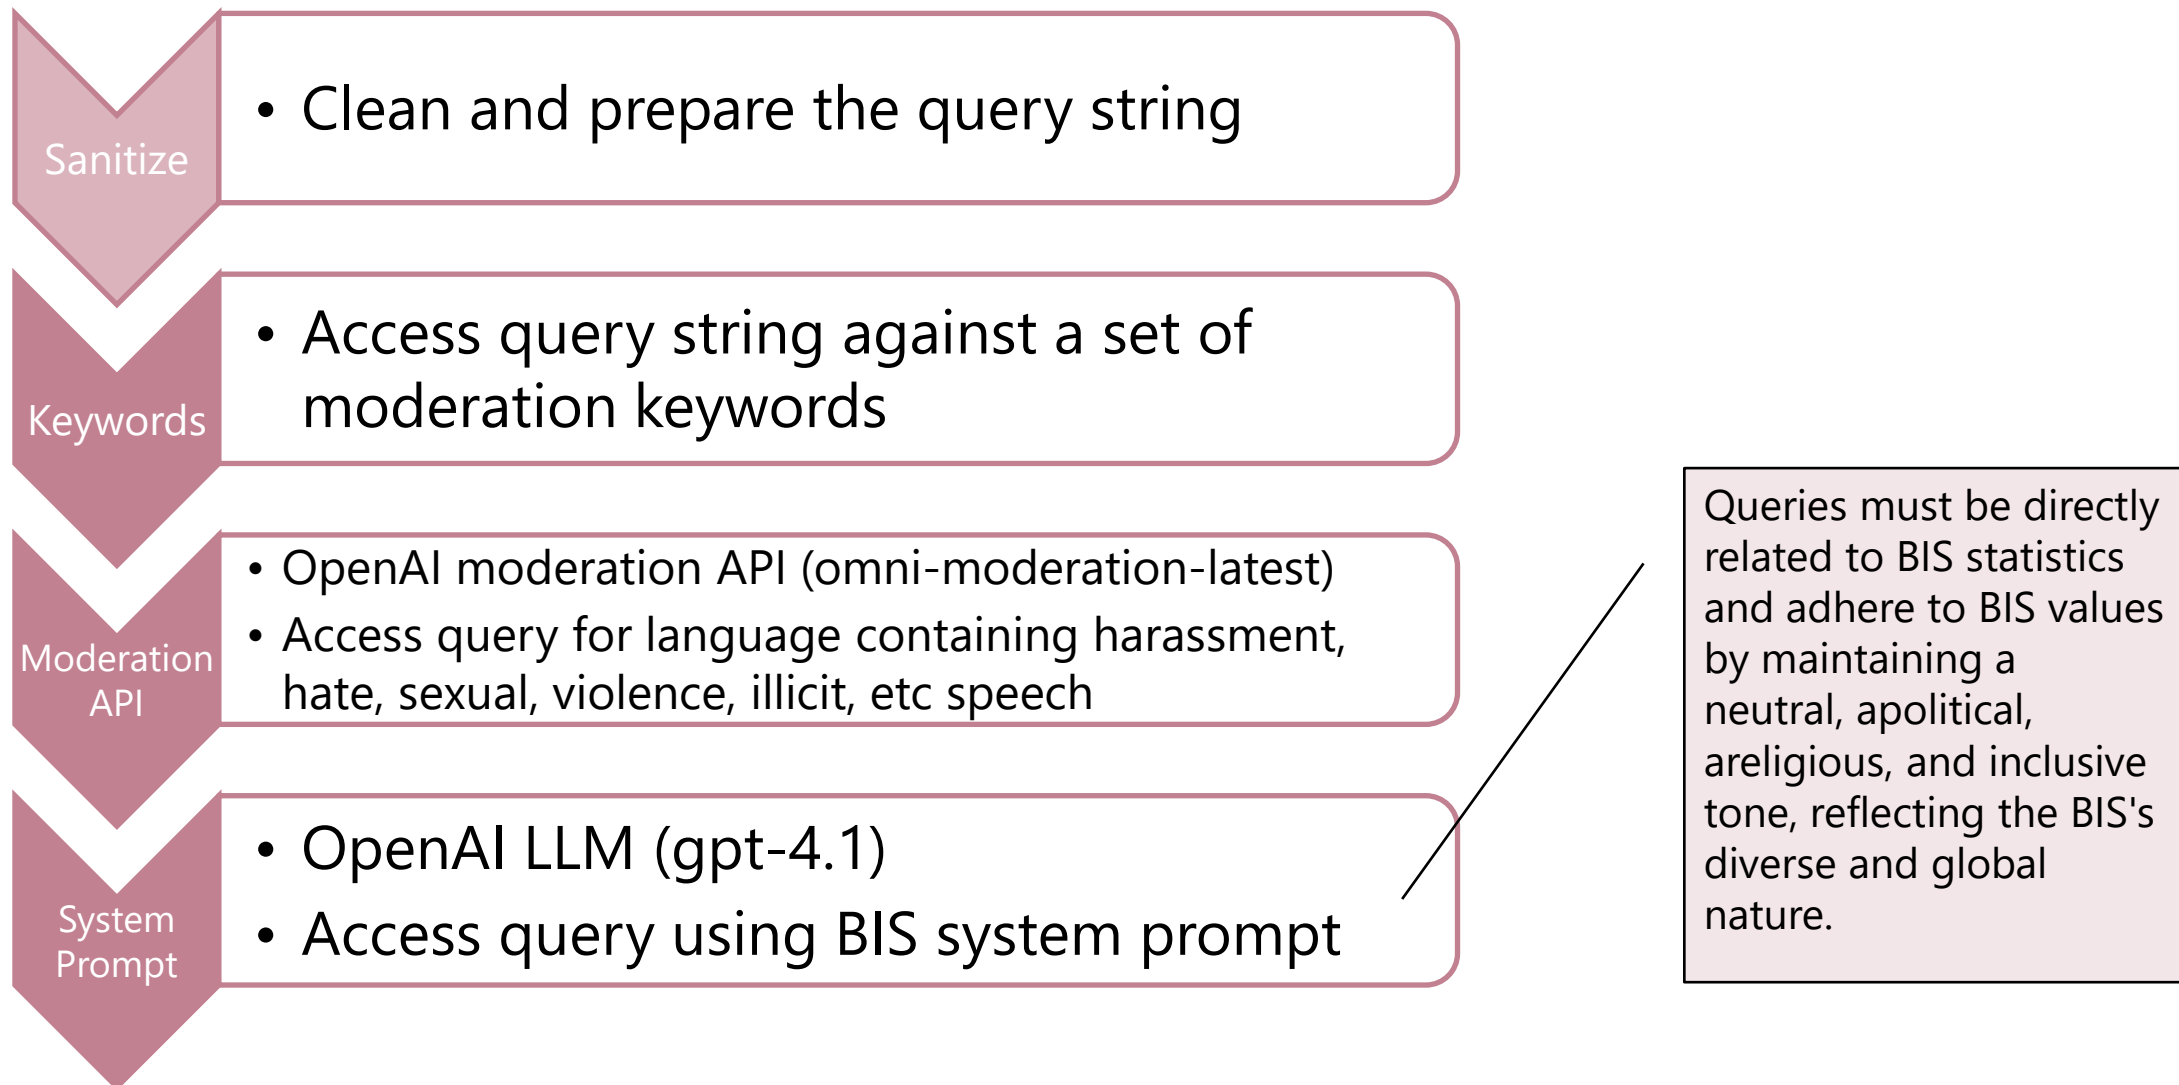
## Annex

## BIS Data Portal – SDMX at its core

Main SDMX artefacts employed:

- **Category Schemes:** Organize and define statistical topics within the data universe.
- **Data Flow:** Represents a BIS-published data collection, outlining what data is available.
- **Dataset Definition (DSD):** Specifies the structure and rules of a published data collection.
- **Concepts:** Define the statistical ideas or terms used in the dataset, dimensions (identifying) or attributes (describing) (e.g. Country, Indicator, Confidentiality).
- **Codelist:** Provides an enumerated list of possible values for each concept (e.g., country codes, indicator types, confidentiality labels).

# AI assisted search - query moderation





# Evaluation

Testing with **OpenAI endpoints**:

- Started with **GPT-4o-mini**
- **GPT-4o** better rationale / results
- **GPT-4.1** cheaper / good results

	Input Context Window	Max Output Tokens	Tokens	Cost
GPT-4o	128k	16,384	1M input tokens	\$2.50
			1M cached input tokens	\$1.25
			1M output tokens	\$10.00
GPT-4.1	1,047,576	32,768	1M input tokens	\$2.00
			1M cached input tokens	\$0.50
			1M output tokens	\$8.00

	Avg Query Time Performance Topic Extraction (sec)	Avg Query Time Performance Country Extraction (sec)	Avg Query Time Performance Dimensions Extraction (sec)
GPT-4o	2,7721	1,3126	2,5452
GPT-4.1	2,1553	1,7330	3,1392

# Lessons learned

## ● Query moderation

- Crucial that the new functionality aligns with the BIS corporate values / positioning
- Layered approach:
  - Keyword filter → OpenAI Moderation → OpenAI System Prompt
- Few-shot examples of acceptable / unacceptable queries improve moderation

## ● AI-search

- Layered approach provides superior results
  - Topic → Reference Area → Dimensions
- Requests to the LLM should be clear and discrete (Input context + System Prompt)



Demo – backup slides



Q What are the office space prices in Rabat

× | ⚙ Explore

⚙ Include AI-powered results ☒


Results for **What are the office space prices in Rabat**

Time series

Include AI-powered results ☒

🌟 AI-powered time series suggestions

Show reasoning ?



Preparing Suggestions...



Results for **What are the office space prices in Rabat**

Time series

Include AI-powered results ☒

✦ AI-powered time series suggestions | [Show reasoning](#) ⓘ

[Commercial property prices in the Reference area Morocco](#)

Did you find the results useful?  



### Show Reasoning

Topic reasoning

Region reasoning

Dimension reasoning

The query is about office space prices in Rabat, which falls under the broader category of commercial property. The most relevant topic in the dictionary is 'Commercial property prices', as it specifically tracks developments of office, retail premises, and industrial property prices.



### Show Reasoning

Topic reasoning

Region reasoning

Dimension reasoning

Rabat is the capital city of Morocco. According to the rules, if a city is mentioned, its corresponding geographic area from the provided list should be returned. Therefore, 'Morocco(MA)' is identified as the relevant geographic area.



### Show Reasoning

Topic reasoning

Region reasoning

Dimension reasoning

The query specifically asks about office space prices in Rabat. Rabat is the capital city of Morocco, so 'Capital city/biggest city/financial center' is selected for 'Covered area'. 'Office & Retail' is selected for 'Real estate type' because the query is about office space. No other dimensions are explicitly mentioned in the query, so their values are left blank.



×

### Why wasn't this useful?

☐ Incomplete

☐ Too generic

☐ Not clear enough

☐ Not what I was looking for

☐ Other

Share some details

Tell us why you gave this feedback

Your feedback is anonymous; please ensure it does not contain any personal or identifying details.

Cancel

Submit

✓

### Thank you for your feedback

×

Your input helps us improve and provide a better experience.