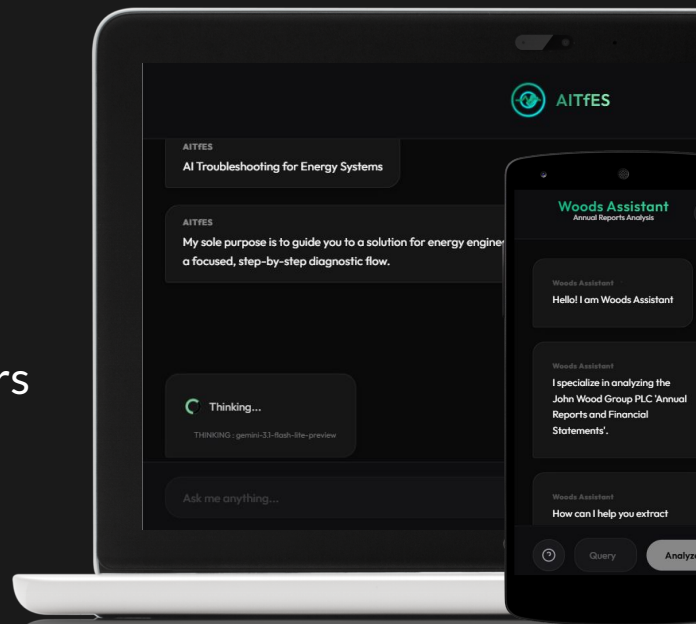




A domain-adaptable modular AI system that delivers context grounded responses using configurable system instructions and RAG pipelines.





What is AITfES ?

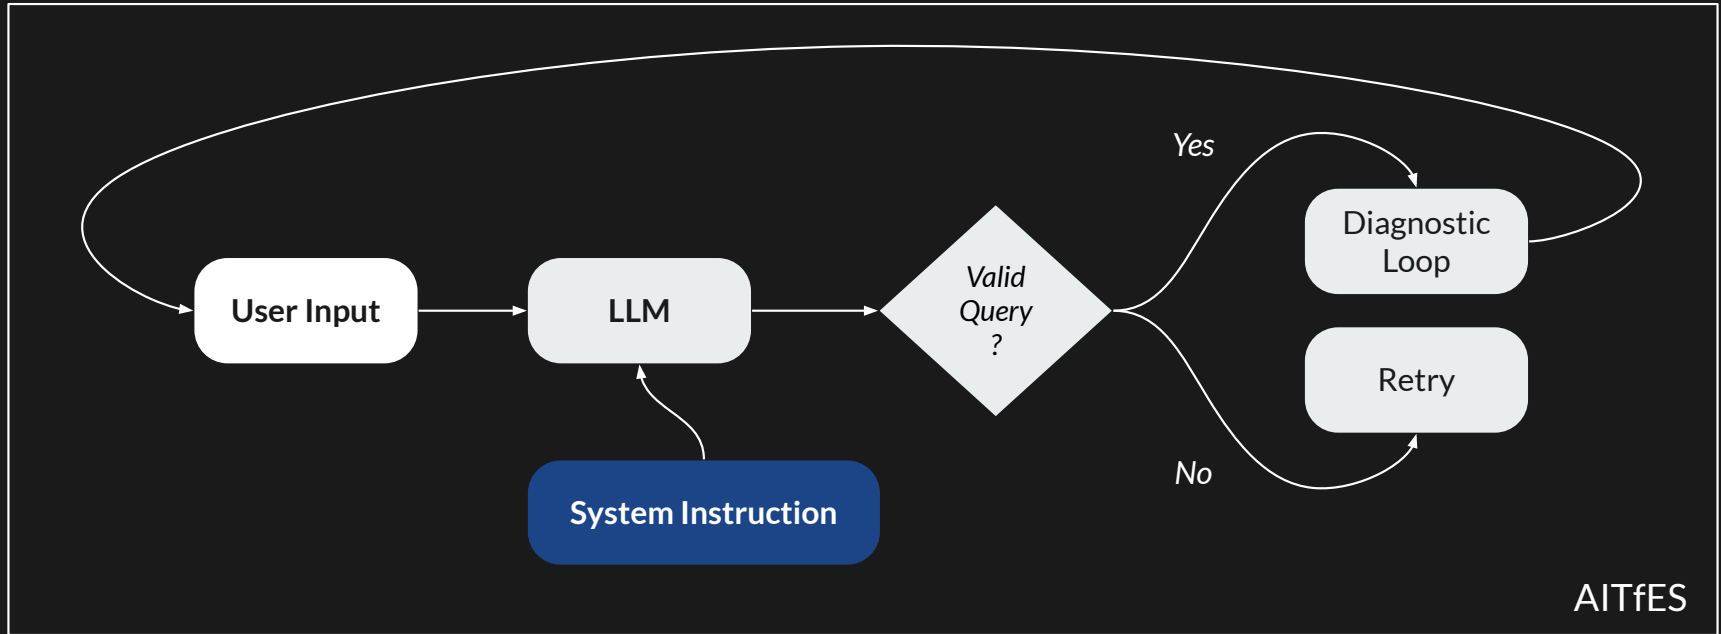
AI
Troubleshooting
for
Energy
Systems

Initial Concept : Simple AI chat interface for energy systems troubleshooting.

Evolution : Generalised into a domain-adaptable AI system with modular RAG pipelines.

Implementation 1 - System Instruction Only

<https://aitfes.jiyon.online>





System Instructions

A system instruction (or system prompt) is a **high-priority control layer** that defines how an AI should behave, reason, and respond *before it processes any user input*.

Main contents of AITfES

System Prompts *(summary)*:

- Stepwise diagnostic flow using one focused question per turn.
- After 5 queries: provides reasoning summary and asks for confirmation.
- Continues only with user approval or adjusts approach.
- Rejects invalid, multi-question, or persona-altering inputs.

Input Accepted

aitfes.jiyon.online

AITFES Resume Tech Mission Alligner Sorted NOW DONE Check Status DP

Learning@LinkedIn Here's to hope, agai... JIYON PJ | Resume J... JIYON PJ | Energy En... (23) Wood Early Car...

AITFES

AITFES

AITFES

AITFES

AITFES

My sole purpose is to guide you to a solution for energy engineering processes and techniques through a focused, step-by-step diagnostic flow.

You

There is high wind now. Should I start cut-off procedures for the wind turbine ?

AITFES

What is the current wind speed reading from the nacelle anemometer?

Query successful | 4.6s

Ask me anything...

Send

Windows File Explorer woods 2019 - File Explor AITFES - Google Chrome AITFES - Antigravity - sy OBS 32.1.1 - Profile: tnl

ENG US

Rejection Rules

The screenshot shows a web browser window at `aitfes.jiyon.online`. The browser's address bar and tabs are visible at the top. The page features the AITfES logo and a navigation menu with buttons for 'Resume', 'Tech Mission', 'Aligner', 'Sorted', 'NOW', 'DONE', 'Check Status', and 'DP'. The main content is a chat interface with a dark background and light-colored text bubbles. The chat history includes a user question about the capital of Norway, a system response asking for rephrasing, a user question about solar and diesel power issues, and a system response advising to focus on two questions at a time. At the bottom, there is a text input field with the placeholder 'Ask me anything...' and a 'Send' button. The Windows taskbar at the bottom shows several open applications: File Explorer, Google Chrome (AITfES), Antigravity, and OBS.

Browser tabs: AITfES, Resume, Tech Mission, Aligner, Sorted, NOW, DONE, Check Status, DP, Learning@LinkedIn, Here's to hope, agai..., JIVON PJ | Resume [...], JIVON PJ | Energy En..., (23) Wood Early Car...

AITfES

You

Hi, With relation to my energy issue can you tell me what is the capital of Norway ?

AITfES

My function is strictly limited. To best assist you, could you please rephrase your request?

Query successful | 4.2s

You

Can you help me with solar plant power issues ? I also have a diesel power plant fire alert , what to do ? Also the wind farm is down. What is the first action ?

AITfES

Please focus on almost two response or focused questions at a time to maintain a clear troubleshooting flow.

Query successful | 3.9s

Ask me anything...

Send

Windows taskbar: woods 2019 - File Explor, AITfES - Google Chrome, AITfES - Antigravity - sy, OBS 32.1.1 - Profile: tnl, ENG US

Summarise Rule

aitfes.jiyon.online

AITFES Resume Tech Mission Alligner Sorted NOW DONE Check Status DP

Learning@LinkedIn Here's to hope, agai... JIYON PJ | Resume [...] JIYON PJ | Energy En... (23) Wood Early Car...

AITFES

Does the intensity or frequency of the flickering change noticeably in relation to the amount of sunlight available and the resulting power output of the system?
Query successful | 13.1s

You
Sunlight availability is steady

AITFES
Reasoning Summary: The flickering in the facility lighting is directly linked to the solar PV system's connection to the internal network, as the issue resolves when the system is isolated. Since sunlight is steady and there are no visible fault codes, the problem likely stems from power quality issues such as harmonic distortion, rapid voltage fluctuations caused by the inverter's switching behavior, or a potential impedance mismatch between the inverter output and the facility's load.
Confirmation Request: Should I proceed with diagnostics focused on the inverter's power quality and the electrical stability of the connection to the facility?
Query successful | 5.7s

Ask me anything... Send

woods 2019 - File Explor AITFES - Google Chrome AITFES - Antigravity - sy OBS 32.1.1 - Profile: tnl

ENG US



Gaps Identified

Broad knowledge, but lacking precision.

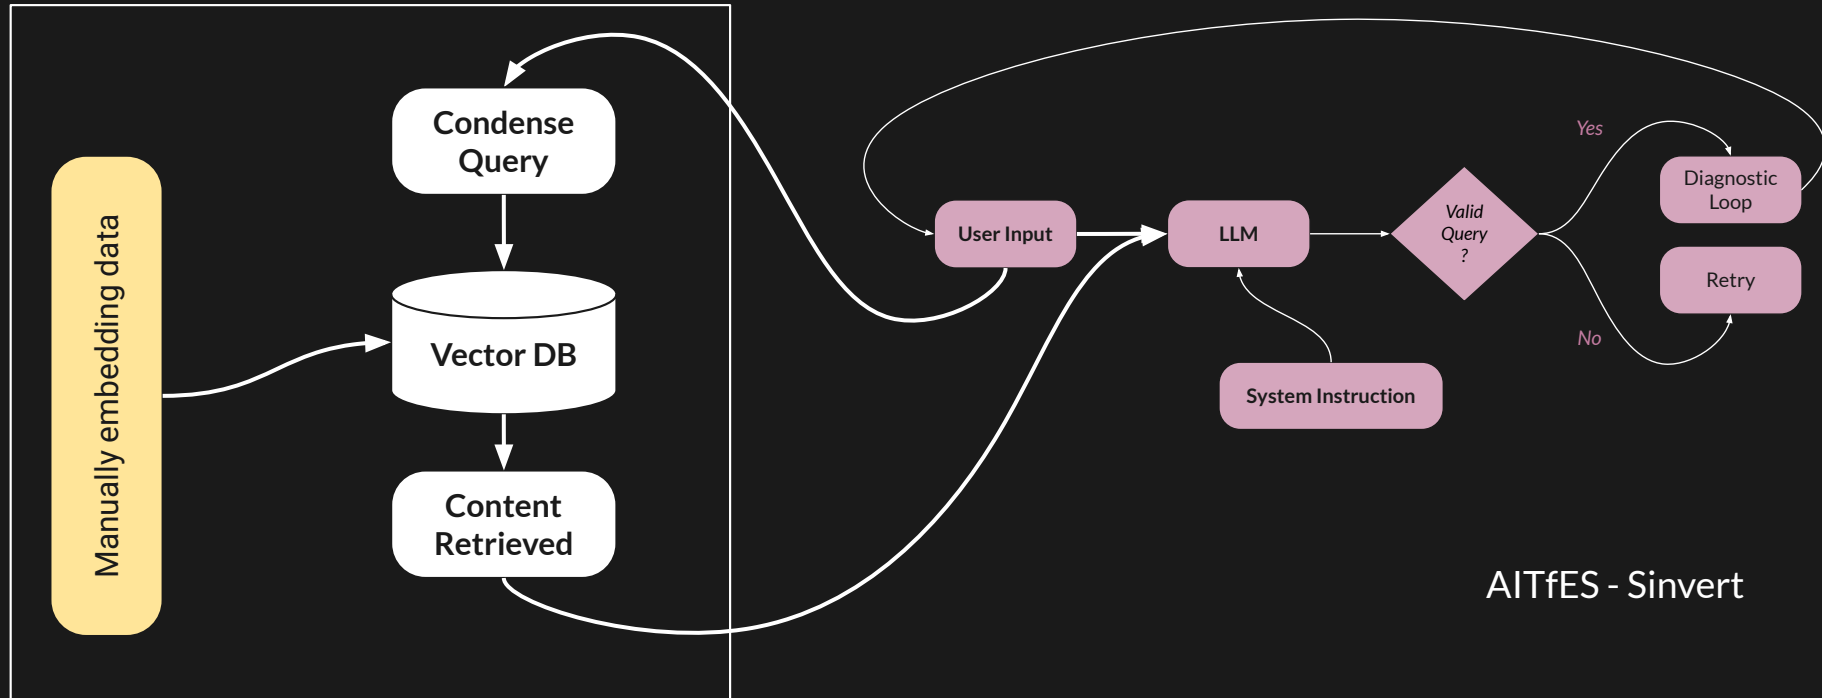
Lacks focus, targeting, and contextual relevance.

Several options were considered to fill the gap and **RAG** seemed most promising.

- Uses relevant context for better answers.
- Reduces hallucinations with grounded data.
- Adds domain focus without retraining.

Implementation 2 - System Instruction + RAG

<https://aitfes.jiyon.online/sinvert>



AITFES - Sinvert



AITfES - Sinvert

SINVERT PVS500–600 series are high-efficiency, transformerless 3-phase central solar inverters by Siemens for large-scale PV plants.

The first document selected to embed into a vector database was,

“SINVERT PVS 500 / PVS 600 - Operating Instructions · 04/2011”

- **Primary Objective was to learn how to implement RAG.**

- Get started
- Database
- Quickstart
- Indexes (2)
- Backups
- Assistant
- Inference
- API keys
- Manage

STARTER USAGE ⓘ

WUs ⓘ 58K / 2M

RUs ⓘ 4 / 1M

Storage ⓘ 0.056GB / 2GB

Upgrade now

sinvert ⓘ

Record count 148 • Region ^{aws} us-east-1 | Model llama-text-embed-v2 | Host sinvert-e11hp62.svc.aped-4627-b74a...

Connect

- BROWSER
- METRICS
- NAMESPACES (1)
- CONFIGURATION

Records

Upsert record

Namespace: sinvert-pvs500-600 | Operation: Search by text | Text: troubleshoot | Top K: 10

+ Filter + Rerank Search

Search: 10 results (top_k=10) 5 embed tokens ⓘ 1 RUs ⓘ

1	_id: rec147 category: "index" text: "Topics include installation, safety, operation, maintenance, troubleshooting, and technical specifications."
SCORE	0.1921


https://aitfes.jiyon.online/sinvert

The screenshot shows a web browser window with the URL `aitfes.jiyon.online/sinvert`. The browser's address bar and tabs are visible at the top. The main content area features the AITFES logo and the text "SINVERT PVS 500 / PVS 600". Below this, there is a chat interface with a header "I'm AITFES 📌 - SINVERT Support". The chat history includes a system message from AITFES: "Specialized diagnostic flow for SINVERT PVS 500 / PVS 600." followed by another system message: "I can search the knowledge base for SINVERT PVS 500 / 600 issues. Ask me anything about these products." A user message in a blue bubble asks: "If my sinvert pvs 500 is overheating, What to do ?". A final system message from AITFES asks: "What specific fault or alarm messages, if any, are displayed on the inverter?" with a status of "Query successful | 20.0s". At the bottom, there is a text input field with the placeholder "Ask me anything..." and a "Send" button.

aitfes.jiyon.online/sinvert

AITFES Resume Tech Mission Alligner Sorted NOW DONE Check Status DP

Learning@LinkedIn Here's to hope, agai... JIYON PJ | Resume J... JIYON PJ | Energy En... (23) Wood Early Car...

 **AITFES**

SINVERT PVS 500 / PVS 600

I'm AITFES 📌 - SINVERT Support

AITFES

Specialized diagnostic flow for SINVERT PVS 500 / PVS 600.

AITFES

I can search the knowledge base for SINVERT PVS 500 / 600 issues. Ask me anything about these products.

You

If my sinvert pvs 500 is overheating, What to do ?


AITFES

What specific fault or alarm messages, if any, are displayed on the inverter?

Query successful | 20.0s

Ask me anything...

Send



Better contextual
relevance, narrow
scope and focused
troubleshooting.

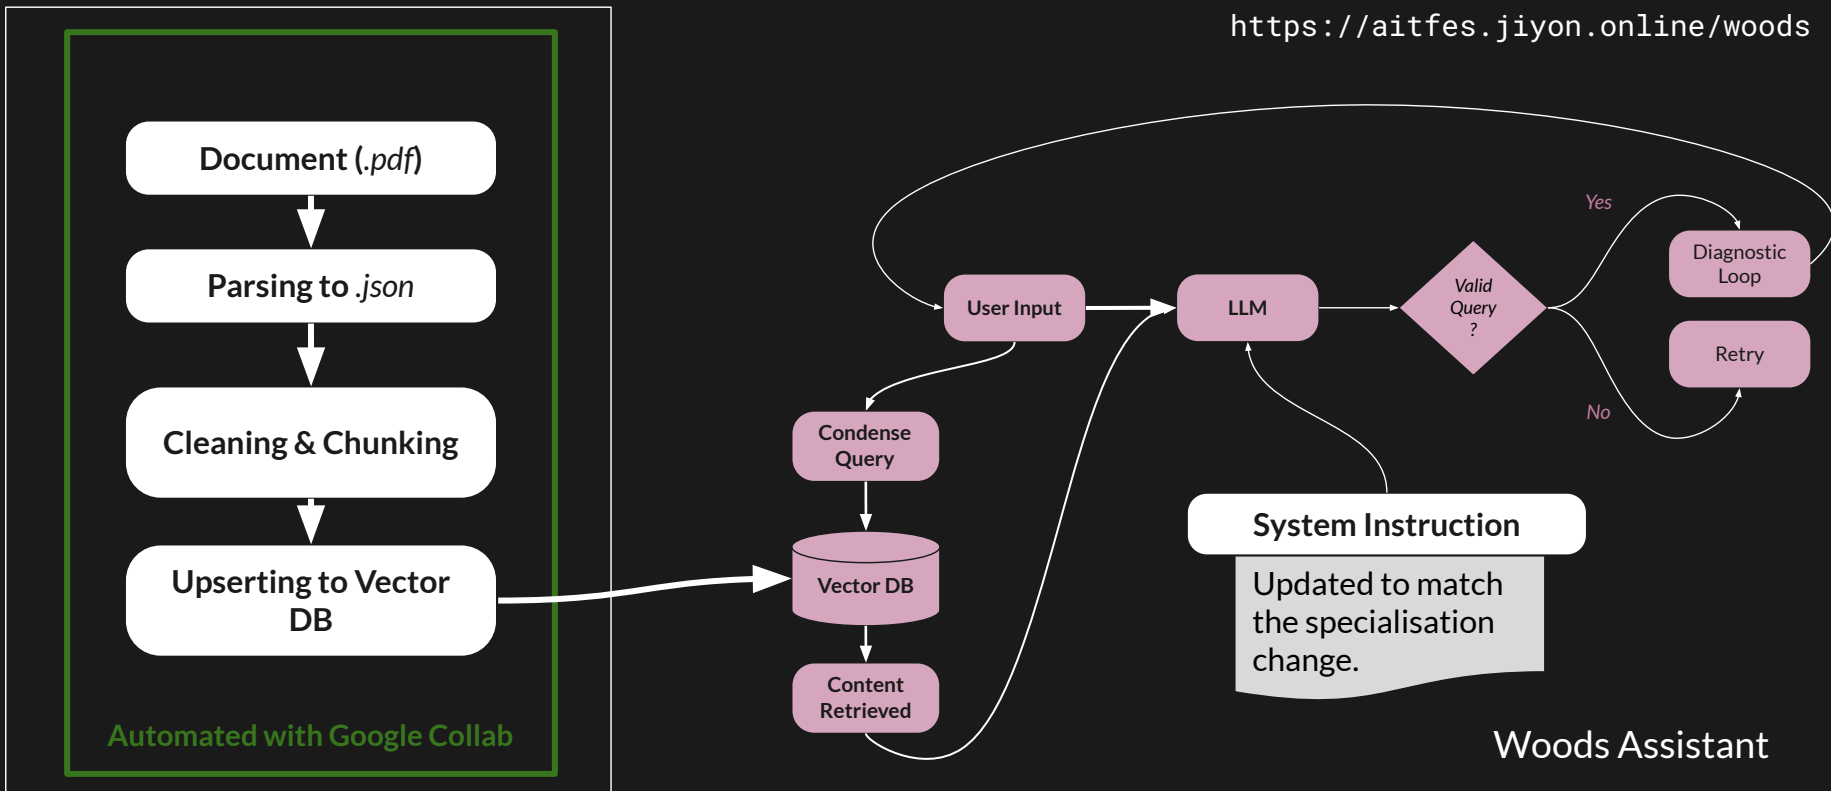
Gaps Identified

Quality and quantity of data is very important.

- Manually converting document to embed ready format is time consuming.
- Low quality data (**148 records, only 1 document**) causes inaccurate retrieval.
- Better data ingestion is needed.

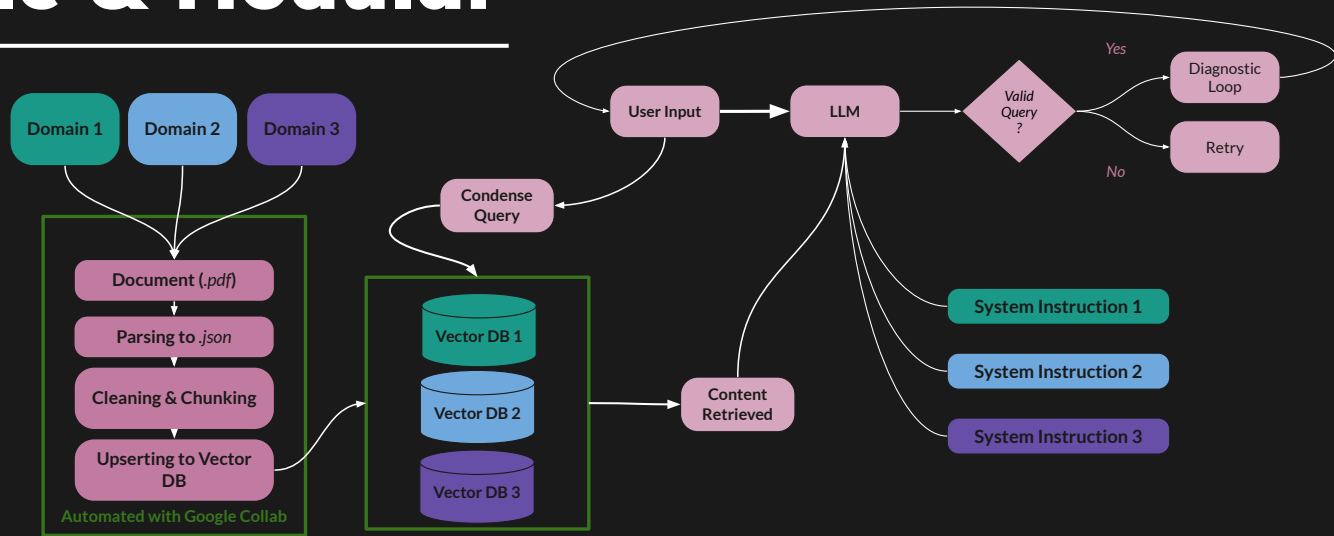
Implementation 3 - System Instruction + RAG + Automated Data Ingestion

<https://aitfes.jiyon.online/woods>



Woods Assistant

Scalable & Modular



Data Scaling: Automated ingestion enables continuous knowledge base expansion.

Modular Architecture: Decoupled LLM, retrieval and application layers allow independent upgrades or component replacement.

Domain Scalability: Same architecture adapts across domains by changing data and system instructions.



Woods Assistant

To establish scalability beyond the initial specialisation, a custom assistant for Wood Annual Reports was set-up.

- *Automated ingestion* of Wood Annual Reports from 2019 - 2024
- Automated data ingestion means more records in the database in minimal time is possible.
- *Modified System Instruction* to reflect the new perspective.

colab.research.google.com/drive/101xX95uKA02IARRPac9QYxwZIHjy

AITFES Resume Tech Mission Aligner Sorted NOW DONE Check Status DP

PDF to Upsert Pipeline - Pinecone [TEMPLATE].ipynb

File Edit View Insert Runtime Tools Help

Commands + Code + Text Run all

PDF to Upsert Pipeline - Pinecone

IMPORTANT

- Do not run all cells at once.
- Each is sequential.
- You must upload files at certain points, check terminal.

JSON file is derived from PDF

- Provide link in field. A link is necessary. - *line 14*
- If link is not available, option to upload PDF will be active.

```
1 # @title
2 !pip install pdfplumber requests tqdm tiktoken
3
4 import pdfplumber
5 import requests
6 from tqdm import tqdm
7 import os
8 import re
9 import ison
```

Variables Terminal

colab.research.google.com/drive/101xX95uKA02IARRPac9QYxwZIHjy

AITFES Resume Tech Mission Aligner Sorted NOW DONE Check Status DP

PDF to Upsert Pipeline - Pinecone [TEMPLATE].ipynb

File Edit View Insert Runtime Tools Help

Commands + Code + Text Run all

```
123
124 print("Upload complete.")
```

Choose files No file chosen Upload widget is only available when

Saving parsed_output.json to parsed_output (7).json

Loaded 421 records

Prepared 1843 chunks for Pinecone

Uploaded batch 1

Uploaded batch 2

Uploaded batch 3

Uploaded batch 4

Uploaded batch 5

Uploaded batch 6

Uploaded batch 7

Uploaded batch 8

Uploaded batch 9

Uploaded batch 10

Uploaded batch 11

Uploaded batch 12

Uploaded batch 13

Uploaded batch 14

Uploaded batch 15

Uploaded batch 16

Uploaded batch 17

Uploaded batch 18

Uploaded batch 19

Uploaded batch 20

Uploaded batch 21

Uploaded batch 22

Uploaded batch 23

Uploaded batch 24

Upload complete.

Variables Terminal

Wood Annual Reports
(2019 - 2024)

Parsing to .json

```
{  
  "page": 35,  
  "type": "text",  
  "content": "for, [content omitted for brevity] 2019 ",  
  "tokens": 394  
}
```

Cleaning & Chunking

```
{  
  "_id": "woods2019_page_35_text_1",  
  "text": "Reduce [content omitted for brevity] data",  
  "category": "audit_and_control_issues"  
}
```

Upserting to
Pinecone

```
{  
  "id": "woods2019_page_35_text_1_chunk_3",  
  "values": [ 0.00919342,  
    .....  
    0.00573730469 ],  
  "metadata": {  
    "category": "audit_and_control_issues",  
    "text": "Due [content omitted for brevity] basis." } }
```

- Get started
- Database
 - Quickstart
 - Indexes (2)
 - Backups
- Assistant
- Inference
- API keys
- Manage

STARTER USAGE ⓘ

WUs ⓘ 58K / 2M

RUs ⓘ 6 / 1M

Storage ⓘ 0.056GB / 2GB

Upgrade now

woods ⓘ

Record count 11,986 • **Region** ^{aws} us-east-1 **Model** llama-text-embed-v2 **Host** woods-e11hp62.svc.aped-4627-b74a...

Connect

BROWSER METRICS NAMESPACES (1) CONFIGURATION

Records

Upsert record

Namespace	Operation	Text	Top K
woods-financial	Search by text	2019	10
+ Filter	+ Rerank		Search

Search: 10 results (top_k=10)

6 embed tokens ⓘ 1 RUs ⓘ

1 **_id:** woods2022_page_60_text_1_chunk_6
category: "financial_information"
text: "In 2019, as part of our strategic planning process, we undertook Opportunities qualitative scenario planning exploring the pace and depth of the low-c..."
SCORE
0.3225

Browser navigation bar: [aitfes.jiyon.online/woods](#)

Navigation tabs: PENS PEN, AITFES, Resume, Tech Mission, Aligner, Sorted, NOW, DONE, Check Status, DP

Woods Assistant

Annual Reports Analysis

Woods Assistant: I am Woods Assistant

Woods Assistant: Specializes in analyzing the John Wood Group PLC 'Annual Reports and Financial Statements'.

Woods Assistant: How can I help you extract data or analyze across different financial periods today?

You: How did the financial performance compare between 2023 and 2024?

Query:

Analyze

Browser navigation bar: [aitfes.jiyon.online/woods](#)

Navigation tabs: PENS PEN, AIT FES, Resume, Tech Mission, Alligner, Sorted, NOW, DONE, Check Status, DP

Woods Assistant

Annual Reports Analysis

Woods Assistant

The financial performance of John Wood Group PLC in 2024 showed a significant decline compared to 2023, characterized by reduced profitability, substantial non-recurring charges, and ongoing challenges in cash generation ([John Wood Group PLC Annual Report and Financial Statements 2024](#)).

Financial Performance Comparison

The following table summarizes the key financial metrics for 2023 and 2024:

Metric	2023 (Restated)	2024
Revenue	\$5,561m	\$5,490m
Adjusted EBIT	\$169m	\$81m
Adjusted EBITDA	\$388.2m	\$284.1m
Operating Loss	\$(55m)	\$(2,631m)
Free Cash Flow	\$(313m)	\$(153m)

Buttons: Query, Analyze



Increasing the amount of high-quality data in the existing RAG system significantly improved its reliability.

- Automating data ingestion led to significant increase in records in the database.
- **148 records** in Implementation 2 vs almost **12,000 records** in Implementation 3
- Better data quality led to better retrieval and LLM accuracy.

	IMPLEMENTATION 1	IMPLEMENTATION 2	IMPLEMENTATION 3
Core Approach	System Instructions driven	System Instruction + RAG	System Instruction + RAG + automated data ingestion
Data Source	None	Single document	Multi-document
Data Scale	No external data	~150 records	~12,000 records
Ingestion Method	-	Manual data ingestion	Automated parsing, chunking and upserting
Capability	LLM Controlled responses	Context-aware answers	Cross-document, domain-adaptable reasoning

Tech Stack

- LLM Layer : **Gemini API** - *access to multiple llm models*
- Application Layer : **SvelteKit** - *ui, server-side logic, routing*
- Retrieval Layer : **Pinecone** - *vector database, data embedding*
- Development : **Antigravity** - *agent-first IDE, code assist*
- Data Pipeline : **Colab** - *data ingestion pipeline*
- Deployment : **Vercel** - *serverless hosting*

Note: Designed and deployed using only free-tier services (*net cost \$0*) for cost-efficient and rapid prototyping.

Conclusions

A **modular AI** system for **context-grounded** responses, designed to **scale across multiple domains**.

- Built and deployed an end-to-end AI system prototype at zero cost.
- Iteratively improved from prompt-based control to scalable RAG architecture.
- Applied to real engineering and financial datasets.

Limitations:

- Performance depends on data quality and coverage.
- No real-time integration (static document-based system).
- Evaluation is qualitative, not benchmarked.





LINKS

All implementations are live for demonstration.

Woods Assistant :

→ aitfes.jiyon.online/woods

AITfES - Sinvert:

→ aitfes.jiyon.online/sinvert

AITfES :

→ aitfes.jiyon.online

Upsert Pipeline :

→ colab.research.google.com/drive/101xX9...

GitHub :

→ github.com/thenewlegend/AITfES

Thank You

