




WORKSHOP:

Ask your Drupal site questions
and get proper answers

Frederik Wouters 06/03/2024



Throwback 🧓



Dialogflow Phone Gateway ^{BETA}


**Here's my number,
So call me maybe
+1 201-365-7047(english only)**

Examples


- Run cron
- Flush cache
- Put site offline/
- weather
- joke



Press `esc` to exit full screen



HTML



**CMS ON
STRUCTURED DATA**



**DYNAMIC
CONTENT-TYPES
AND FIELDS**



**AI
FILLING DRUPAL
CONTENT**






3 Minutes of Theory







Parts of a RAG

- Embeddings
- Vector Database
- LLM (Large Language Model)





Vector Database







Large Language Model

Deep learning algorithm to generate human sounding text

Important concepts

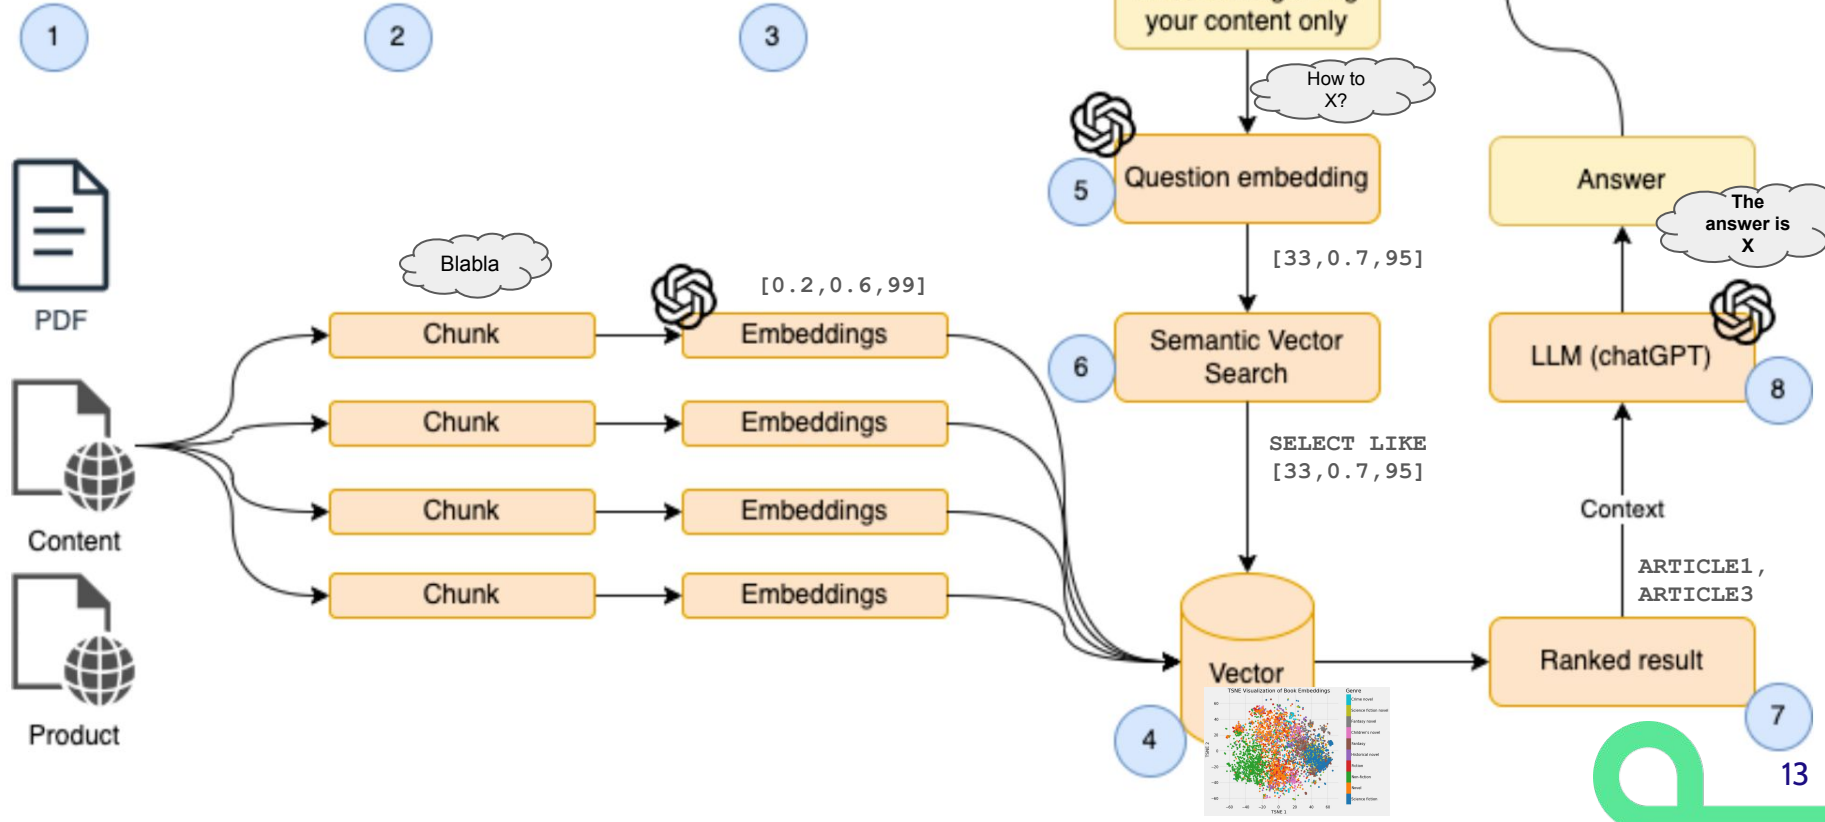
- Prompt
- Tokens
- Temperature





When these come together:

Magic happens




LET'S GO





Starting up gitpod/Drupal



1 Select

woutersf/workshop-mountain-camp-ch-24

New Workspace

Create a new workspace in the woutersf organization.

<> Select a repository

> VS Code · 1.87.0
Editor · Browser

🔧 Standard
Class · Up to 4 cores, 8GB RAM, 30GB storage

Continue


2 Continue



```
▼ CONTAINERS
  ▼ workshop-mountain-camp-ch-24
    > ▶ mailhog - Up Le
    > ▶ phpmy - pal9_project_pma
    > ▶ wodby - 9_project_php - U
    > ▶ wodby - 9_project_crond -
    > ▶ wodby - pal9_project_mar
    > ▶ wodby - 19_project_nginx -
```

Compose Logs
Compose Start
Compose Stop
Compose Restart
Compose Down


1 Stop

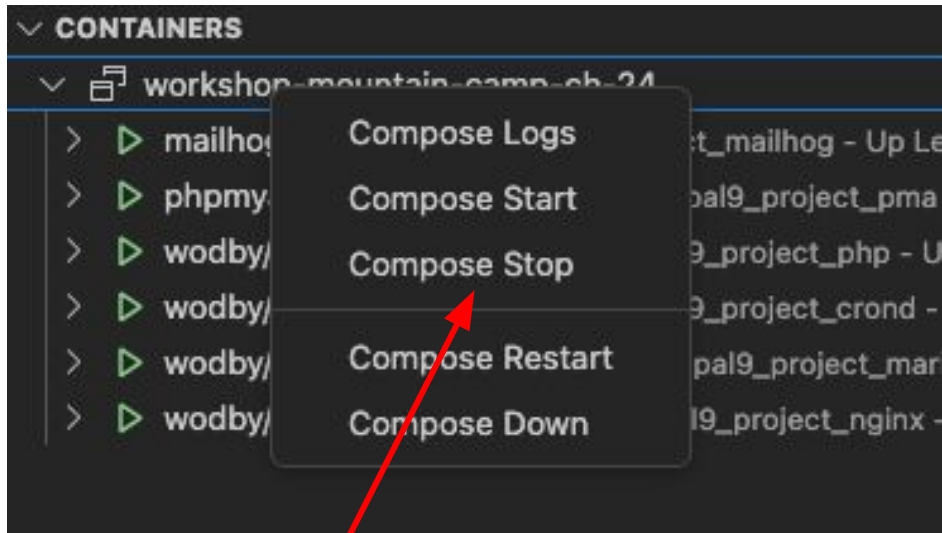




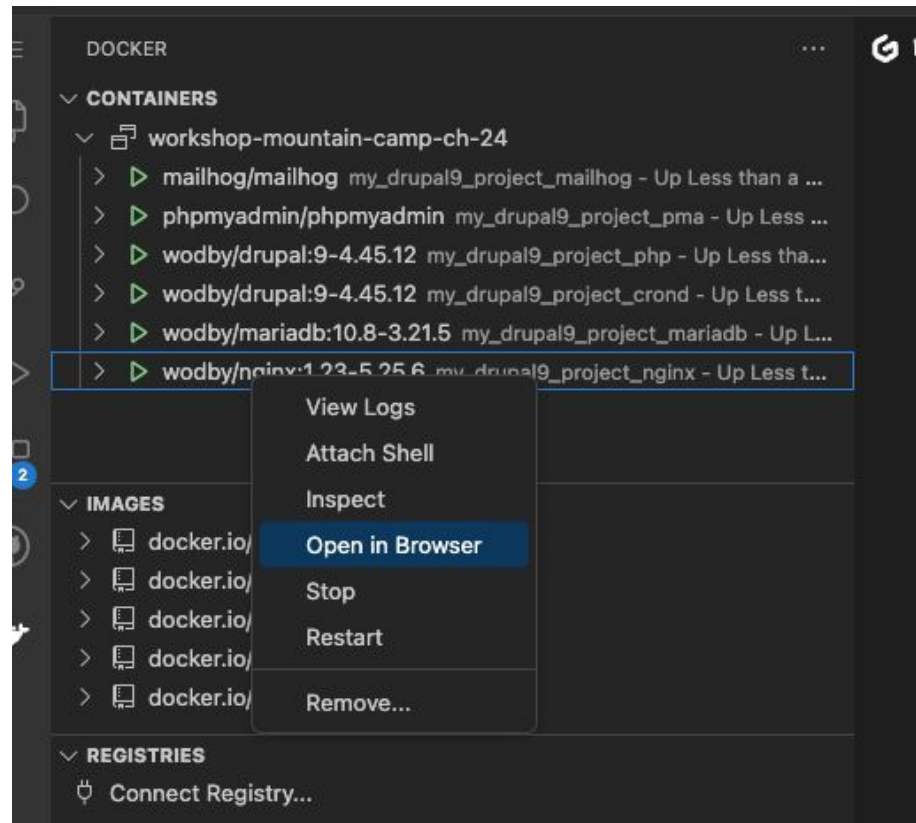
Before you start clicking.

```
Sudo composer install
```





1 Start





Install Drupal

The screenshot shows the Drupal 10.2.3 installation wizard. The browser address bar at the top displays the URL: 88-woutersf-workshopmounta-aliaicj8yt.ws-eu108.gitpod.io/core/install.php. The main content area features a vertical list of six steps on the left, with the first step, '1 Choose language', highlighted in a light blue bar. To the right of this list, the heading 'Choose language' is displayed above a dropdown menu currently set to 'English'. Below the dropdown is a blue button labeled 'Save and continue'.

88-woutersf-workshopmounta-aliaicj8yt.ws-eu108.gitpod.io/core/install.php


Drupal 10.2.3

- 1 Choose language
- 2 Choose profile
- 3 Verify requirements
- 4 Set up database
- 5 Install site
- 6 Configure site

Choose language

English

Save and continue



Errors

^ Errors found

✘ File system

Writable (*public* download method)
The directory `sites/default/files` does not exist. An automated attempt to create this directory failed, possibly due to a permissions problem. To proceed with the installation, either create the directory and modify its permissions manually or ensure that the installer has the permissions to create it automatically. For more information, see `INSTALL.txt` or the [online handbook](#).

✘ Settings file

The *Settings file* does not exist. The Drupal installer requires that you create a `./sites/default/settings.php` as part of the installation process. Copy the `./sites/default/default.settings.php` file to `./sites/default/settings.php`. More details about installing Drupal are available in [INSTALL.txt](#).

(see WORKSHOP.md)

```
docker exec -it my_drupal10_project_php /bin/bash

sudo mkdir web/sites/default/files

sudo chmod 777 web/sites/default/files

sudo cp web/sites/default/default.settings.php /
web/sites/default/settings.php

sudo chmod 777 web/sites/default/settings.php
```

Database configuration

Database type*

MySQL, MariaDB, Percona Server, or equivalent

PostgreSQL

SQLite

Database name

Database username

Database password

▼ Advanced options

Save and continue

(From .env file)

Database

- drupal

Username

- drupal

Password

- drupal

ADVANCED!!!

Host

- mariadb



Congratulations, you installed Drupal!



Welcome!


You haven't created any frontpage content yet.

Congratulations and welcome to the Drupal community.

Drupal is an open source platform for building amazing digital experiences. It's made, used, taught, documented, and marketed by the [Drupal community](#). Our community is made up of people from around the world with a shared set of [values](#),




Enabling & configuring





Enable the following modules

- Admin toolbar
- search_api
- Search API AI
- Search API Pinecone
- OpenAI ChatGPT Explorer






Configure Openai Module

- `/admin/config/openai/settings`
(find also in `WORKSHOP.md`)

See DRUPAL SLACK #AI for keys

(Hint: Save it twice)





TEST Openai Configuration

- `/admin/config/openai/chatgpt`
(find also in `WORKSHOP.md`)

OpenAI ChatGPT Explorer ☆

Ask ChatGPT *

write short oneliner joke about a proprietary CMS

Enter your text here. When submitted, OpenAI will generate a response from its Chats endpoint. Please note that each query counts against your API usage for OpenAI. Based on the complexity of your text, OpenAI traffic, and other factors, a response can sometimes take up to 10-15 seconds to complete. Please allow the operation to finish. Be cautious not to exceed the requests per minute quota (20/Minute by default), or you may be temporarily blocked.

Options

Response from OpenAI

Why did the proprietary CMS refuse to go to the open-source party? Because it didn't want to share its code on the dance floor! 😊 Do you have a favorite CMS you prefer to use?

The response from OpenAI will appear in the textbox above.

Sorry annoying bugfix

```
sudo vi  
web/modules/contrib/openai/modules/openai_embeddings/src/Form/SettingsForm.php
```

Comment out validate handler contents

(find also in [WORKSHOP.md](#))

```
    */  
    public function validateForm(array &$form, FormStateInterface $form_state) {  
        /*parent::validateForm($form, $form_state);  
  
        foreach ($this->pluginManager->getDefinitions() as $pid => $plugin) {  
            $plugin_instance = $this->pluginManager->createInstance($pid);  
            if (!method_exists($plugin_instance, 'validateConfigurationForm')) {  
                continue;  
            }  
            $subform = &$form['connections'][$pid];  
            $subform_state = SubformState::createForSubform($form['connections'][$pid],  
                $plugin_instance->validateConfigurationForm($subform, $subform_state);  
        }  
    }  
    */}
```

Configure embeddings

- `/admin/config/openai/openai-embeddings/settings`
(find also in **WORKSHOP.md**)

See DRUPAL SLACK #AI fo

(Hint: Save it twice)

Configure API clients for vector search database services

Vector database login*

Milvus
✓ Pinecone

Choose the vector database to store embeddings in.

^ **Pinecone**


Configure Pinecone settings.

API Key*

The API key is required to make calls to Pinecone for vector searching.



Search_api 🔥





Configure Search API


Configure Search API

1 Add Server

2 Add index

3 Add Fields (title / body)

4 field type embeddings



^ Configure *Pinecone* backend

Namespace

An optional override for the default namespace. The index ID will always be appended unless namespaces are disabled in the P configuration.

Save



Your own cool namespace

^ Bundles

Which bundles should be indexed?

Only those selected

All except those selected

Bundles

Article

Basic page

Server

- No server -

Pinecone server

Select the server this index should use. Indexes cannot be ena

Save and add fields

Add Article our test CT



Select Pinecone



Add fields to index

^ Content

Label

Machine name

Property path

Body

body

body

Title

title


title


Add Title and body field



Switch fields to embedding

Set to **embeddings**
for textfields to index





Try to save a node

(add in some wikipedia content on a interesting subject).

88-woutersf-workshopmounta-aliaicj8yt.ws-eu108.gitpod.io/admin/co... ☆

Back to site Manage Shortcuts admin Announcements

Content Structure Appearance Extend Configuration People Reports Help

View Edit Fields Processors

Index status




1/1 indexed

100%



Chat block



Add Block to Content region

The screenshot shows the Drupal 8 admin interface for managing blocks. The browser address bar shows the URL: `88-woutersf-workshopmounta-aliaccj8yt.ws-eu108.gitpod.io/admin/structure/block`. The top navigation bar includes "Back to site", "Manage", "Shortcuts", and "admin". A secondary navigation bar contains icons for "Content", "Structure", "Appearance", "Extend", "Configuration", "People", "Reports", and "Help".

The main content area displays a list of blocks with their placement options. A red arrow points to the "Content" block placement option in the "Content" region.

Block Name	Region	Placement Options
Social Bar	Social Bar	Place block
RSS feed	System	Social Bar
Content Above	Content Above	Place block
Page title	core	Content Above
Help	Help	Content Above
Content	Content	Place block
Main page content	System	Content

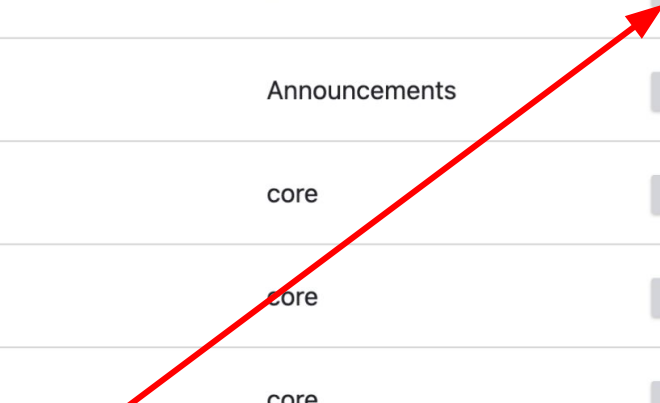


Add Block to Content region

Place block

Filter by block name

Block	Category	Operations
Search API AI: Chat form	AI	Place block
Announcements Feed	Announcements	Place block
Page title	core	Place block
Primary admin actions	core	Place block
Tab	core	Place block



^ Chat configuration

Chat model

gpt-4 (currently beta invite only) ▾

Select which model to use to analyze text. See the [model overview](#) for details about each model.

Temperature

0,4

What sampling temperature to use, between 0 and 2. Higher values like 0.8 will make the output more random, while lower values like 0.2 make it more focused and deterministic.

Max tokens

1024



Configure Block

System message

You are a chat bot to help find resources and provide links and references from the User's private knowledgebase. You will base all your answers off the provided context that you find from the user's knowledgebase. Always return links as HTML.

This is the first message sent to the LLM and sets the overall tone of how the LLM will respond. Note LLMs prioritize User messages over System messages.

Assistant message

I found the following information in the User's Knowledge Base:
[context]

I will respond with information from the User's Knowledge Base above.

Top K

The number of results to return for a vector query.

Score threshold

The threshold vector embeddings must meet to be considered relevant. A score of 0,5 means everything should be considered relevant and a threshold of 0,9 means only the most relevant results should be returned.

Maximum query length

The maximum length of the query used to create embeddings.

Model

Save Block positioning.

Content

Place block

↕ * Search API AI: Chat form

AI

Content

Configure

↕ Main page content

System

Content


Configure

Save blocks





Test it out!



Site name

Home



My account

Log out

By [admin](#), 5 March, 2024

test content

test content

Search API AI: Chat form

Ask me a question

Send



Some examples



Jouw gids voor de overheid

Stel je vraag en vind je weg...

 Vertel, wat wil je weten?



Wat weet je over verhuizen naar Nederland?

Antwoord

Als u overweegt om naar Nederland te verhuizen, zijn er verschillende zaken die u moet regelen afhankelijk van uw situatie:

Nederlandse of EU-nationaliteit:

- Als u de **Nederlandse nationaliteit** of de nationaliteit van een ander EU-land heeft, kunt u zonder verdere voorwaarden naar Nederland verhuizen en hier wonen en werken^[1].

Britse nationaliteit:

- Voor Britse staatsburgers die niet ook de Nederlandse nationaliteit hebben, is de situatie veranderd sinds het Verenigd Koninkrijk de EU heeft verlaten. Tot en met 31 december 2020 kon men zonder voorwaarden naar Nederland verhuizen. Na deze datum is een **verblijfsdocument** nodig, waarvoor u aan bepaalde voorwaarden moet voldoen. Deze voorwaarden kunt u vinden op de website van de Immigratie- en Naturalisatiedien

Let op! Het antwoord is door een AI gegenereerd. Raadpleeg altijd de onderstaande bronnen.

5 handige links

gebaseerd op jouw vraag

Rijksoverheid

1. Ik wil van het VK terugverhuizen naar Nederland. Wat moet ik regelen?

Rijksoverheid

2. Wat moet ik regelen als ik in Nederland kom wonen?

Rijksoverheid

Rijksoverheid

Vertel, wat wil je weten?

Formatted response

References to CMS

ZüriCityGPT

Ich bin ZüriCityGPT und ich weiss (fast) alles, was auf stadt-zuerich.ch publiziert ist.

Frag mich etwas über die Verwaltung der Stadt Zürich.

Frage stellen →



Wir sind weiterhin fleissig daran, ZüriCityGPT zu verbessern, zu erweitern und die Daten regelmässig zu aktualisieren. Wir freuen uns darum über jedes Feedback. Oder gar einem Weitererzählen, falls du zufrieden bist mit ZüriCityGPT.

Und Versuche doch auch die Vorabversion vom mehr chat-artigen ZüriCityGPT: chat.zuericitygpt.ch (noch sehr im Experimentierstadium)

Hinweise und Einschränkungen



Kontakt und Feedback



ZüriCityGPT

Ich bin ZüriCityGPT und ich weiss (fast)
alles, was auf stadt-zuerich.ch publiziert ist.

Generated Response

References to CMS

Make me a two day city trip, but I'm not very mobile and the weather is bad

Frage stellen →

To plan a two-day city trip in Zurich with limited mobility and bad weather, you can consider the following options based on the provided articles:

Day 1:

- Visit the museums in Zurich to admire ancient and modern art ([Link 1](#)).
- Explore the Technorama or the Umwelt Arena to learn about the laws of nature ([Link 1](#)).
- Experience a guided tour to discover interesting places in the city ([Link 1](#)).

Day 2:

- Enjoy a day at the Alpamare, a sports and leisure center ([Link 1](#)).
- Visit the Zoo Zurich to get up close with animals ([Link 1](#)).
- Attend the "Zwei Tage Zeit" festival for improvised music at Theater Rigiblick ([Link 4](#)).

Please note that these suggestions are based on the available information and may not cover all possible options. It is recommended to check the specific opening hours and accessibility of each location before planning the trip.



But what about my **DATA**





OS Embeddings

MTEB or Massive Text Embedding Benchmark


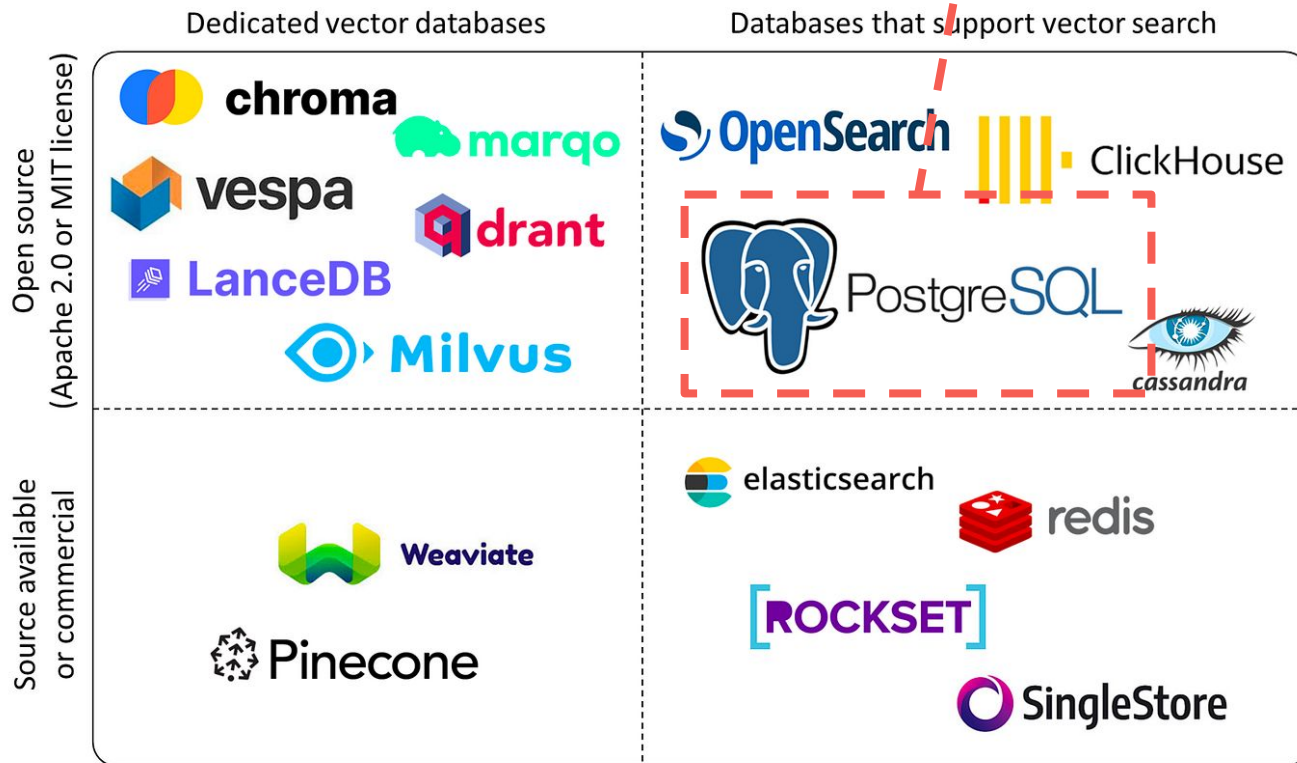
On par with GPT
WITHOUT THE RISKS

Rank ▲	Model ▲	Average ▲	ArguAna ▲	ClimateFEVER ▲	CQADupstackRetrieval ▲	DBPedia ▲	FEVER ▲	FiQA2018 ▲	HotpotQA ▲	MSMARCO ▲	NFCor
1	SFR-Embedding-Mistral	59	67.17	36.41	46.49	49.06	89.35	60.4	77.02	43.41	41.88
2	GritLM-7B	57.41	63.24	30.91	49.42	46.6	82.74	59.95	79.4	41.96	40.89
3	e5-mistral-7b-instruct	56.89	61.88	38.35	42.97	48.89	87.84	56.59	75.72	43.06	38.62
4	voyage-lite-02-instruct	56.6	70.28	31.95	46.2	39.79	91.35	52.51	75.51	37.93	43.7
5	voyage-lite-01-instruct	55.58	58.73	37.47	45.11	43.42	89.71	44.79	70.46	39.66	43.33
6	text-embedding-3-large	55.44	58.05	30.27	47.54	44.76	87.94	55	71.58	40.24	42.07
7	GritLM-8x7B	55.09	59.49	28.69	47.63	46.54	85.02	49.89	73.83	35.55	39.05
8	Cohere-embed-english-v3.0	55	61.52	38.43	41.53	43.36	88.97	42.19	70.72	42.93	38.57
9	UAE-Large-V1	54.66	66.15	38.53	40.93	44.89	88.24	44.84	73.13	41.4	38.65
10	bge-large-en-v1.5	54.29	63.54	36.57	42.23	44.11	87.18	45.02	74.1	42.49	38.13
11	Cohere-embed-multilingual-v3.0	53.84	55.11	29.96	40.64	41	88.53	44.1	70.61	43.45	36.42
12	GIST-large-Embedding-v0	53.44	63.38	33.99	43.44	42.96	86.55	44.3	70.46	41.39	38.65



Vector Database

Run it yourself






OS Large Language Model


Llama 2 is just as accurate as GPT-4 and much faster and more efficient than GPT-3.5 and GPT-4

Some examples

- Bert (proprietary)
- Claude (proprietary)
- GPT-3 (Openai = proprietary)
- GPT-3.5 (Openai = proprietary)
- GPT-4 (Openai = proprietary)
- **Llama 2 (Open Source)**
- **Bloom (Open Source)**



Frederik wouters 90%



Frederik Wouters

Link me on linkedin
[/in/woutersfrederik/](https://www.linkedin.com/in/woutersfrederik/)

+32 498 53 55 76

woutersf@gmail.com