

#### 18.06.2024

### Look mom – I got a LLM in my database!

Maximillian Jahn @maxjahn

# About

#### MMag. Maximillian Jahn, LL.B.(WU)

- Head of Cloud Advisory at Orbit Cloud Solutions
- Former Principal Cloud Architect at Oracle
- Certified Cloud Architect on OCI, AWS, Azure and Google Cloud
- Focus on statistics in university
- Co-author of several scientific papers on recommender systems

### Agenda

Introduction How AI impacts organizations today

**Fundamentals** What's all the buzz about

Using LLMs in the Database Make your Oracle database understand content

The Scotch Badger Shop A hypothetical case study

**Benefits of in-DB LLMs** Best of two worlds

Conclusion Outlook



### Introduction

How AI impacts organizations today

### **Expectations** are high



74% of leaders expect GenAl to impact their business significantly.

https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html

### **Expectations** are high



Q: What are the key benefits you hope to achieve through your generative AI efforts? (Oct./Dec. 2023) N (Total) = 2,835 91% of all organizations expect their productivity to increase due to generative AI.

Now decides next: Insights from the leading edge of generative Al adoption. Deloitte's State of Generative Al in the Enterprise. January 2024

#### Look mom – I got a LLM in my database!

### **Expectations are high**



31%

of the leaders we surveyed expect substantial transformation in less than one year; 48% expect it in one to three years.

> Now decides next: Insights from the leading edge of generative AI adoption. Deloitte's State of Generative AI in the Enterprise. January 2024

### **Concerns are big**



#### Q. What are or will be the top 3 barriers to the implementation of AI techniques within your organization? Source: 2023 Gartner AI in the Enterprise Survey 810529 C

#### Gartner

https://www.gartner.com/en/newsroom/press-releases/2024-05-07-gartner-survey-finds-generative-ai-is-now-the-most-frequently-deployed-ai-solution-in-organizations

### **Concerns are big**

Now: Key findings

#### Preparedness for generative AI

Respondents claimed the highest levels of **preparation in technology and strategy**, while feeling far less prepared in risk and talent.



Figure 7

Q: Consider the following areas. For each, rate your organization's level of preparedness with respect to broadly adopting generative AI tools / applications? (Oct/Dec. 2023) N (Total) = 2,835

Now decides next: Insights from the leading edge of generative AI adoption. Deloitte's State of Generative AI in the Enterprise. January 2024

#### Look mom – I got a LLM in my database!

### **Concerns are big**

Specific generative AI risks and concerns include inaccurate results and information (i.e., "hallucinations"); legal risks such as plagiarism, copyright infringement, and liability for errors; privacy and data ownership challenges; lack of transparency, explainability and accountability; and systemic bias. The latter exemplifies another category of risk in which AI amplifies and exacerbates a problem that already exists, such as propagating and systematizing existing social biases, facilitating and accelerating the spread of misinformation, helping criminals commit crimes, or fanning the flames of political divisiveness.

According to the business and technology leaders we surveyed during fourth quarter 2023, the biggest concerns related to governance were: lack of confidence in results (36%), intellectual property issues (35%), misuse of client or customer data (34%), ability to comply with regulations (33%), and lack of explainability / transparency (31%).

Now decides next: Insights from the leading edge of generative AI adoption. Deloitte's State of Generative AI in the Enterprise. January 2024

### Expectations

- Productivity Increase
- Redefine the work of leaders as much as employees
- GenAl will transform transformation
- Improve existing products and services
- Encourage innovation and growth
- Uncover new ideas and insights

### Concerns

- Estimating and demonstrating Al value
- Lack of talent/skills
- Lack of confidence in results
- Intellectual property issues
- Misuse of client or customer data
- Ability to comply with regulations
- Lack of explainability / transparency



### **Fundamentals**

What's all the buzz about

#### "Any sufficiently advanced technology is indistinguishable from magic."

Arthur C. Clarke's 3rd Law

### Large Language Models

- A large language model (LLM) is a computational model notable for its ability to achieve general-purpose language generation and other natural language processing tasks such as classification.
- Based on language models, LLMs acquire these abilities by learning statistical relationships from vast amounts of text during a computationally intensive self-supervised and semi-supervised training process.
- Coming from the transformer architecture, there are two big groups of LLMs:
  - o Autoencoding Models
  - o Autoregressive Models

### **Autoencoding Models**





Text to be encoded

Statistical relations from training data



Vector representation of input in the semantic space "Embeddings"

### Understand

- Distill from text
  - o Semantic Meaning
  - o Semantic Similarity
  - o Summaries
  - o Categorization
- "What piece is missing?"

• Encoder part of Transformer



#### Visualizations using UMAP (Uniform Manifold Approximation and Projection for Dimension Reduction)



### **Autoregressive Models**



#### Generate

- Generate Tokens
  - o Generate Text
  - o Chat
- "... tell me more"

• Decoder part of Transformer





### Using LLMs in the Database

Make your Oracle database understand content

### Which type to use in a DB?

- In a database, we most likely want to **enrich our understanding of the data**, so we want to leverage autoencoding features.
  - Semantic Similarity
  - Semantic Search
  - o Summarization
  - Categorization

• To get a semantic representation of our data we use **embedding models** 

### **Generating Embeddings**



### **Embeddings Models Available**

- Proprietary Models
  - o OpenAl
  - o Cohere
  - Google AI / VertexAI
  - Anthropic
  - o Mistral

- Open Source Models
  - BERT (Bidirectional Encoder Representations from Transformers)
  - RoBERTa (Robustly Optimized BERT Pretraining Approach)
  - o DistilBERT
  - ELMo (Embeddings from Language Models)
  - Many more at HuggingFace Hub!



Proprietary ("Closed Source") models require endpoints deployed by vendor or partners (e.g. Azure and OpenAl, Oracle and Cohere)

Open Source models can be hosted on own infrastructure

Model	Dimensions	Maximum Tokens	Languages	Size	Parameters		
sentence-transformers/all-mpnet-base-v2	768	512	English	438 MB	110M	$\overline{)}$	
sentence-transformers/all-MiniLM-L6-v2	384	256	English	90.9 MB	22M		
ProsusAl/finbert	768	512	English	417 MB	110M		
medicalai/ClinicalBERT	768	512	English	417 MB	110M		
sentence-transformers/distiluse-base-multilingual-cased-v2	512	512	Multilingual (various)	135 MB	134M		
BAAI/bge-small-en-v1.5	384	512	English	42 MB	33M		
BAAI/bge-base-en-v1.5	768	512	English	90 MB	109M		
taylorAI/bge-micro-v2	128	256	English	25 MB	17M	> Op	p
prajjwal1/bert-tiny	128	128	English	14 MB	4M		
TaylorAl/gte-tiny	128	128	English	25 MB	23M		
intfloat/multilingual-e5-base	512	512	Multilingual (various)	90 MB	278M		
intfloat/multilingual-e5-small	256	256	Multilingual (various)	45 MB	118M		
avsolatorio/GIST-Embedding-v0	768	512	English	438 MB	109M		
Alibaba-NLP/gte-base-en-v1.5	768	8192	English		137M		
Snowflake/snowflake-arctic-embed-m	768	512	English	230 MB	109M	)	
Cohere/embed-english-v3.0	1024	512	English, Multilingual			$\sum$	
OpenAl/text-embedding-ada-002	1536	8191	English	-	1.2B		
OpenAl/text-embedding-3-small	1536	8191	Multilingual	-	-		0
OpenAl/text-embedding-3-large	3072	8191	Multilingual	-	-		
Google Al/google-gecko.text-embedding-preview-0409	768	2048	Multilingual	-	1.2B	$\mathcal{I}$	

24

### **Embedding Model Preferences**

• DB machines usually do not have (or support) GPUs, so the model should be small and lean

- (Open Source) Models can be run directly inside the Oracle DB
  - Oracle Machine Learning (OML) supports in-DB models for some time now

### **Generating Embeddings**



### **Importing Models**



### **HuggingFace Pretrained Models**

- sentence-transformers/all-mpnet-base-v2
- sentence-transformers/all-MiniLM-L6-v2
- sentence-transformers/multi-qa-MiniLM-L6-cos-v1
- ProsusAl/finbert
- medicalai/ClinicalBERT
- sentence-transformers/distiluse-base-multilingualcased-v2
- sentence-transformers/all-MiniLM-L12-v2
- BAAI/bge-small-en-v1.5
- BAAI/bge-base-en-v1.5
- taylorAl/bge-micro-v2
- intfloat/e5-small-v2

- intfloat/e5-base-v2
- prajjwal1/bert-tiny
- thenlper/gte-base
- thenlper/gte-small
- TaylorAl/gte-tiny
- infgrad/stella-base-en-v2
- sentence-transformers/paraphrase-multilingualmpnet-base-v2
- intfloat/multilingual-e5-base
- intfloat/multilingual-e5-small
- sentence-transformers/stsb-xlm-r-multilingual

### **Convert Models to ONNX Format**

- ONNX is an open format built to represent machine learning models.
   ONNX defines a common set of operators the building blocks of machine learning and deep learning models and a common file format to enable AI developers to use models with a variety of frameworks, tools, runtimes, and compilers
- OML4Py 2.0 supports the conversion of models from HuggingFace to ONNX ready for import into Oracle 23ai
- Docker image for easy conversion: maxjahn/oml4py:latest
- Upload to OCI Object Storage and import using PL/SQL

### **Create Embeddings in the DB**

- Initial Load
  - With SQL or PL/SQL
- Automated updates
  - o Database Triggers
  - o DB Scheduler Jobs

All is done within the database: No data is exported or passed to remote API call!

### **Query the Data**

- Add vector indices
  - In-Memory Neighbor Graph Vector Index: Hierarchical Navigable Small World (HNSW)
     → not supported in RAC (= OCI Autonomous DB) environments ☺
  - Neighbor Partition Vector Index: Inverted File Flat (IVF)
- Use Oracle SQL vector functions (VECTOR\_DISTANCE) to query data
  - Euclidian, Cosine, Dot, Manhattan



### The Scotch Badger Shop

A hypothetical case study

### **Case Study Setting**

- Distributor for Scotch whisky
- Usually has ~300 SKU for sale
- Reviews for ~2300 types of whisky
- Stock management in Oracle DB
  - Autonomous DB in OCI with Oracle 23ai
- Apex applications for sales and backoffice
- Customer-facing Website



### **Case Study Requirements**

- Add a **semantic search** to the backoffice application to find products
- Add functionality to find similar products
- Add functionality to find **similar products that are in stock**
- Add a chatbot on the website that can talk about Whisky and give recommendations based on internal documentation and reviews that are already stored in the corporate Oracle database

• Be able to train/fine-tune the models used

### **Model Selection**

- Embedding model
- English language
- No specific domain knowledge
- Fast and small

- For the implementation sentence-transformers/all-mpnet-base-v2 was used
  - Converted to ONNX and uploaded to OCI Object Storage
  - Imported ONNX model with PL/SQL

### **Generate Embeddings**



### **Backoffice Application**

≡ Scotch Badger Backend					
Products	E Reset	+ Create			
Q Search	No Record Selected				
☐ In Stock					
Aberlour 16 year old, 40% Single Malt Scotch					
Adelphi (distilled at Bowmore) 12 year old 2000, 5 Single Malt Scotch					
Ardbeg Auriverdes, 49.9% Single Malt Scotch					
Ardbeg Grooves, 46% Single Malt Scotch					
Ardbeg Supernova (2010 Polosco) - 60.1%					

 $\beta$  vector  $\sim$ 

#### **Products**

Q	Search.

🗌 In Stock

Aberlour 16 year old, 40% Single Malt Scotch

Adelphi (distilled at Bowmore) 12 year old 2000, 5... Single Malt Scotch

Ardbeg Auriverdes, 49.9% Single Malt Scotch

Ardbeg Grooves, 46% Single Malt Scotch

Ardbeg Supernova (2010 Release), 60.1% Single Malt Scotch

Ardbeg Supernova 2015 Release, 54.3% Single Malt Scotch Semantic Search

Semantic Search

seaweed and smoke

#### Search Results

No

Release 1.



Х

+ Create  $e \in e_{\alpha}$  Semantic Search

≡ Scotch Badger Backend		A vector				
Products		중 Reset + Create 론, Semantic Searc				
् talisker						
In Stock	Reviews	🖒 Edit Similar Similar In Stock				
	Name	Talisker 10 year old, 45.8%				
Double Barrel Talisker/Craigellachie, 46% Blended Malt Scotch Whisky	Category	Single Malt Scotch				
	Rating	90				
Douglas Laing Premier Barrel (distilled at Taliske Single Malt Scotch	DescriptionFresh, vibrant. Appetizingly briny, with notes of seaweed, fishnets, lemongrass, smoke, and damp peat. A bed of light toffee provides contrast and balance. It finishes with a knockout punch: powerfully peppery and salty. This whisky remains a beautifully dynamic and bold whisky.					
Douglas Laing Provenance	Stock					
Single Malt Scotch	No stock available					
Douglas Laing Provenance Young & Feisty (distilled Single Malt Scotch						
Talisker 10 year old, 45.8% Single Malt Scotch						
Talisker 10 year old, 45.8% Single Malt Scotch	Release 1.0					

Scotch Badger Backend

 $\beta$  vector  $\sim$ 

#### Similar $\times$ **Products** ₿. Semantic Search + Create Name Category Description Rating Fresh, vibrant. Appetizingly briny, with notes of seaweed, fishnets, Q talisker Single lemongrass, smoke, and damp peat. A bed of light toffee provides Talisker 10 year Malt 90 contrast and balance. It finishes with a knockout punch: old, 45.8% C Edit R powerfully peppery and salty. This whisky remains a beautifully Similar Similar In Stock Scotch In Stock dynamic and bold whisky. Antique gold color. Aromas of peat smoke, damp earth, exotic **Double Barrel** pepper, vanilla, and fruit. Medium body. A great balance of flavors, Talisker/Craigellachie, 46% with a lovely Islay character. The whisky begins sweet up front, Blended Malt Scotch Whisky ultimately turning pleasingly dry towards the finish. In between Single you'll enjoy flavors of peat smoke, vanilla, spicy pepper, and 89 Malt Cadenhead lemongrass, smoke, and damp peat. A bed of **Douglas Laing Premier** delicately honeyed malt, with just a hint of the sea. What I like Scotch out punch: powerfully peppery and salty. This Barrel (distilled at Taliske... most about this whisky is its impeccable balance-between youth Single Malt Scotch and maturity, sweetness and dryness-with just the right amount of peat smoke to declare its Islav roots without smothering its other flavors. **Douglas Laing Provenance** (distilled at Talisker) 6... Straw-gold color. On the nose, sweet toffee, citrus notes, seaweed, Single Malt Scotch and spice complement a powerful peat smoke infusion. In body, it is thick and oily. On the palate, a somewhat sweet maltiness up **Douglas Laing Provenance** front is run over by a powerful peat smoke locomotive. Again, the Young & Feisty (distilled... Single whisky is enriched with citrus and pear notes, spice, and seaweed. Ardbeg, 10 year Single Malt Scotch Malt 93 The finish is powerful, long, and warming. The smoke lingers for old. 46% Scotch minutes, if not hours. If you like your Ardbeg to go to a phenolic extreme, you will cherish this one. This big, powerful whisky Talisker 10 year old. 45.8% makes no apologies for its Islay roots. And the fact that this Single Malt Scotch whisky is bottled at 46% ABV just makes this big whisky even bigger. Talisker 10 year old, 45.8% Release 1 ~ Single Malt Scotch Light to medium in body (and quite pale in color), but quite

≡ Scotch Badger Backend		Q vector
Products		중 Reset + Create 邑, Semantic Sea
् talisker		
In Stock	Reviews	🖒 Edit Similar Similar In Stock
	Name	Talisker 10 year old, 45.8%
Double Barrel Talisker/Craigellachie, 46% Blended Malt Scotch Whisky	Category	Single Malt Scotch
	Rating	90
Douglas Laing Premier Barrel (distilled at Taliske Single Malt Scotch	Description	Fresh, vibrant. Appetizingly briny, with notes of seaweed, fishnets, lemongrass, smoke, and damp peat. A bed of light toffee provides contrast and balance. It finishes with a knockout punch: powerfully peppery and salty. This whisky remains a beautifully dynamic and bold whisky.
Douglas Laing Provenance (distilled at Talisker) 6	Stock	
Single Malt Scotch	No stock available	
Douglas Laing Provenance Young & Feisty (distilled Single Malt Scotch		
Talisker 10 year old, 45.8% Single Malt Scotch		
Talisker 10 year old, 45.8% Single Malt Scotch	Release 1.0	

≡ Scotch Badger Backend

 $\beta$  vector  $\sim$ 

 $\mathbf{E}_{\mathbf{q}}$  Semantic Search

 $\times$ 

#### Similar in Stock

Products					·		+ Create	E. Semantic Se
	Name	Category	Rating	Description		- create		
C       talisker         In Stock         Double Barrel         Talisker/Craigellachie, 46%         Blended Malt Scotch Whisky	R N C R	Ardbeg, 10 year old, 46%	Single Malt Scotch	93	Straw-gold color. On the nose, sweet toffee, citrus notes, seaweed, and spice complement a powerful peat smoke infusion. In body, it is thick and oily. On the palate, a somewhat sweet maltiness up front is run over by a powerful peat smoke locomotive. Again, the whisky is enriched with citrus and pear notes, spice, and seaweed. The finish is powerful, long, and warming. The smoke lingers for minutes, if not hours. If you like your Ardbeg to go to a phenolic extreme, you will cherish this one. This big, powerful whisky makes no apologies for its Islay roots. And the fact that this whisky is bottled at 46% ABV just makes this big whisky even bigger.	Edit	Similar	Similar In Stock
Barrel (distilled at Taliske Single Malt Scotch		Signatory (distilled at Caol Ila), 14 year old, 1990 vintage, 56.2%	Single Malt Scotch	85	Light to medium in body (and quite pale in color), but quite invigorating. A flirtatious vanilla sweetness at the front of the	out punch: p	powerfully per	opery and salty. This
Douglas Laing Provenance (distilled at Talisker) 6 Single Malt Scotch	S				whisky evolves further, expressing notes of smoke, tar, chewy seaweed, brine, feint green tea and olives. A persistent, if delicate, honey sweetness hangs on throughout to offer some gentility. Lingering, smoky finish. A raw whisky, but somehow compelling.			
Douglas Laing Provenance Young & Feisty (distilled Single Malt Scotch				An excellent value too! (Bottled exclusively for Binny's Beverage Depot.)				
		Lagavulin, 12 year, 57.8%	Single Malt Scotch	84	Style: Islay single malt scotch Color: Pale straw Aroma: Peat smoke-lots of it. Damp earth. Hint of olives and seaweed. Palate:			
Talisker 10 year old, 45.8% Single Malt Scotch	Release 1.0				Surprisingly gentle in texture and on the light side of medium in body. A malty sweetness balances the dry-ish peat smoke. Again the damp earth mid palate. An appetizing saltiness emerges			
Talisker 10 year old, 45.8% Single Malt Scotch					towards the finish, but is fleeting. Smoky finish.			
					Very dry, fragrant and spicy. Medium bodied, but firm. Initial,			

### **Enhanced Similarity Search**

- With multimodel search in Oracle, we can query both vector data and relational data to answer questions like
  - Which similar products are in stock or will be restocked next week?
  - Which of our products are very different from our 3 top selling products?
  - Get customers who ordered products that go well with BBQ in the last 6 months? (To send out some promotional offer)
  - Do our customers in Waldviertel have different preferences than those from Vorarlberg?
- These types of queries do not work well with separate vector databases or separate semantic search engines



### Chatbot

### **Retrieval Augmented Generation**

- Combines semantic search with text generation
- Review data and details for >2200 whiskies come from the DB, Oracle Autonomous Database is used as vector store.
- The results are then included in a prompt to an autoregressive (Chat) model
  - Open AI GPT-4, GPT-3.5, GPT-3.5-turbo (external endpoint)
  - Cohere Command (running on Oracle OCI GenAI)
- Chainlit used for the Chatbot UI





I'm looking for some new cocktail recipes to try that are based on Scotch Whisky. Can you recommend a few that are easy to make at home?

#### Used Retriever

Ah, you're after some fine Scotch whisky cocktails, are you? Well, you're in luck my friend. Here are a few favourites that'll have you feeling like you're supping at the heart of Scotland from the comfort of your own home:

#### 1. Rob Roy

This one's a classic like a good old tale, and it's as simple as can be. You'll need:

- 60ml Scotch whisky
- 20ml sweet vermouth
- Dash of Angostura bitters Stir these together with ice, strain into a glass, and garnish with a cherry. Like a Manhattan, but with a Scottish twist.

#### 2. Blood and Sand

This one's a wee bit more complicated, but it's worth it for the show. You'll need:

- 20ml Scotch whisky
- 20ml orange juice
  20ml sweet vermouth



• 20ml cherry brandy Shake 'em all up with ice, then strain into a glass. Garnish

#### 



#### ſĊ

**(**37)

#### which scotch whisky can you recommend for the rob roy?

#### 💰 Used Retriever

Ah, for a Rob Roy, you're wanting a Scotch whisky that can stand up to the sweet vermouth and bitters. Here are a few that would do the job quite handsomely:

- Johnnie Walker Black Label 12 year old, 40%: This blended Scotch has a rich mix of toffee, caramel, almond, and banana flavours, all wrapped up in a cloak of white smoke. It's well-balanced and has a complexity that'll add depth to your Rob Roy.
- Bowmore, 18 year old, 43%: A single malt from Islay, this whisky brings a sherried sweetness that'll complement the sweet vermouth wonderfully. It's got notes of toffee, ripe fruit, tar, dark chocolate, cocoa powder, and a hint of sea salt.
- Smoking Ember, 40%: This one's for the peat lovers. It's big and smoky, with a rich, malty underbelly. There's also tar, licorice, seaweed, and a hint of ginger in the mix. It'll give your Rob Roy a robust, well-rounded character.
- 4. Black Bottle, 10 year old, 43%: A unique blend of seven different Islay whiskies, this one brings a rounded Islay character to the table. It's got peat smoke, brine, seaweed, and a firm, malty foundation.
- 5. Dalmore 1995 The Distillery Manag Exclusive, 40%: This single malt from Dalmore is a bit of a treat. It's got a nose of worn leather candied orange malt, and
- : Type your message here...

V



XX)

-<u>o</u>-

ලි

۲

2

Some runs have been hidden. <u>Show 9 hidden runs</u>

🙆 ChatOpenAl

#### un Feedback Metadata

HUMAN

which scotch whisky can you recommend for the rob roy?

#### HUMAN

\*Whisky Name: Johnnie Walker Black Label 12 year old, 40%\*\nNeed you ask? An unmistakable classic and an exemplar of a blended Scotch whisky, famous around the world. Toffee, swirling caramel, whole almond, and mashed banana amid twisting white smoke. The palate has such poise, balance, and dexterity that it sets the standard for many less accomplished blends to aspire to. Conspicuously iconic, and the closest thing you will find to a complete whisky at this price.

Playground

Add to Dataset

\*Whisky Name: Bowmore, 18 year old, 43%\*\nThis replaces the 17 year old in the standard Bowmore portfolio. (The 17 year old will now be Duty Free Only.) The difference is significant. This one is more sherried than the 17 year old, and it wears the sherry well. I know that many of you, like myself, adore the 17 year old. Indeed, I feel that the 17 year old displays the true essence of Bowmore. But this 18 year old integrates the sherry notes with the Bowmore signature to create a richly flavored Islay whisky experience. Thick toffee, ripe fruit, and tar-tinged smoke are accentuated by dark chocolate, cocoa powder, burnt raisin, dates, and sea salt. Long, entertaining finish.

\*Whisky Name: Smoking Ember, 40%\*\nBig and smoky, but with a rich, balancing, malty underbelly. Tarry rope, leafy bonfire, rooty peat, burnt licorice, smoked seaweed, coal ash, and vanilla-tinged malt, all with emerging brine, subtle Spanish olive, and a hint of ginger. A robust, yet well-rounded, Islay-style whisky with a pleasing oily texture, though I do

#### HUMAN

{'session\_id': 'fd78c327-a1e7-49b5-b11b-f3d8d9e5f682', 'input': 'which scotch whisky can you recommend for the rob roy?'}

ି Run ID ି Trace ID

Annotate

DEVELOPER

#### tell me more about the first recommendation

#### 病 Used Retriever

3

#### Regarding the Johnnie Walker Black Label 12 year old:

This is a rich, complex blend of whisky from the four corners of Scotland. It's a true icon, renowned for its balanced flavour and distinctive smoky character. This blend is created using only whiskies aged for a minimum of 12 years from the four corners of Scotland, Johnnie Walker Black Label has an unmistakably smooth, deep character.

It's made with about 40 different whiskies, each aged at least 12 years. The flavour is quite full and complex, with notes of vanilla, peat smoke and a sweet, creamy caramel. You might also detect a faint fruitiness, which comes from the grain whiskies in the blend.

This is a whisky that's perfect for sipping on its own, but its rich flavours also shine through in cocktails like the Rob Roy. The smoky, peaty elements of the Johnnie Walker Black Label provide a nice balance to the sweetness of the vermouth, creating a cocktail with depth and character.

So, if you're looking for a whisky to create a Rob Roy with a punch, the Johnnie Walker Black Label 12 Year Old might just be the dram you're after, lad! And if you're looking to purchase a bottle, you might want to cherce ut Scotch Badger's online store at https://www.scotchbadger.maxian.io/ for some selection of whisky.

∃ Type your message here...





# Benefits of in-DB LLMs

Best of two worlds

#### **Benefits**

#### • Data Proximity and Performance

- When the LLM is integrated directly into the database, it eliminates the need for data to be extracted and moved between systems. This results in significantly improved performance and reduced latency.
- With the LLM co-located with the data, queries and analyses run faster, providing quicker results and enabling faster decision-making.
- Data both from vectors and relational tables can be used for queries and processing seamlessly together.

#### Scalability and Efficiency

• In-database processing reduces the overall computational resources required, leading to cost savings and improved energy efficiency.

#### **Benefits**

#### • Real-time Analytics

- In-database deployment enables real-time analytics and decision-making. As data is ingested or updated, the LLM can immediately process it and provide insights.
- This capability is especially valuable for time-sensitive decisions, fraud detection, or responsive customer service scenarios.

#### Simplified Data Management

- With the LLM inside the database, data management becomes more straightforward. There is no need for complex data pipelines or separate data preparation steps, as the LLM has direct access to the source data.
- This simplifies the maintenance and updates of the data, ensuring the LLM always works with the most current and accurate information.

#### **Benefits**

#### • Security and Compliance

- In-database deployment enhances security and simplifies compliance. By processing data queries within the secure environment of the database, sensitive information is not exposed outside the database perimeter.
- This deployment method adheres to regulatory requirements, especially when dealing with personally identifiable information (PII) or industry-specific data privacy standards.

#### Granular Access Controls

- Databases typically offer robust access control mechanisms, which can be extended to the LLM. This allows for fine-grained control over who can interact with the LLM and which data they can access.
- This level of control ensures that data remains secure and compliant while enabling collaboration and controlled data sharing.

#### **Use Cases**

#### **Enhanced Customer Profiling**

Combine customer demographics (relational data) with purchase behavior to create a comprehensive customer profile.

#### **Customer 360 and Sentiment Analysis**

Relational tables store customer demographics, purchase history, and interaction records, while vector store holds embeddings of customer feedback, reviews, or social media posts. Deploying an LLM in the database enables sentiment analysis, allowing businesses to understand customer opinions and feelings expressed in natural language.

#### **Contextual Advertising**

LLMs can generate embeddings for usergenerated content. By combining this vector data with user profile information stored in relational tables, a sensitive and contextaware advertising system can be built.

#### **Personalized Product Recommendations**

By embedding product descriptions and user preferences using an LLM and storing the embeddings in a vector store within the same database, the system can provide personalized product recommendations.

#### **Use Cases**

#### **Fraud Detection**

By embedding transaction descriptions and customer behavior patterns using an LLM, the system can detect anomalies and potential fraud.

The LLM can process and understand textual reports, case notes, and regulatory documents, linking them to relevant data points in the relational tables.

#### Job Candidate Matching

By embedding the textual data using an LLM and storing the embeddings in a vector store, the system can enable efficient candidate-job matching.

#### **Clinical Decision Support in Healthcare**

By embedding medical literature, research papers, and clinical guidelines using an LLM and storing the embeddings in a vector store, the system can provide evidence-based recommendations to healthcare professionals

#### Content Recommendation and Personalization

A content platform can store articles, videos, and user interaction data in relational tables. By embedding the content and user preferences using an LLM and storing the embeddings in a vector store, the system can provide personalized content recommendations.

#### **Use Cases**

Obviously, all typical use cases for Retrieval Augmented Generation are relevant too.

please write a short instagram about johnnie walker black label. make it fun and engaging and stay within your personality.

#### 🐧 Used Retriever

Ahoy there, whisky wanderers! Gather round as we set sail with a wee dram of none other than the illustrious Johnnie Walker Black Label. Legend has it, this amber treasure is a blend of about 40 whiskies, each aged for at least 12 years. Makes you feel like you're cradling a wee bit of history in your glass, doesn't it, lads and lasses?

Now, let's dive nose first into this dram. Och, the rich waft of smoke, a hint of vanilla sweetness, oh, and a wee whisper of that peaty charm we all love. It's like a bonfire night in the Highlands. Now take a sip and let's journey together across the palate - a swirl of creamy toffee, a sprinkle of spice, and yes, there's that smoky kiss again.

Now, don't rush off just yet, there's more to this tale. The finish, my friends, is one for the books. It's long, warming and as comforting as a woolen kilt on a frosty Scottish night. So, here's to Johnnie Walker Black Label, a true stalwart in our whisky world. Slainte!

Remember to sip responsibly, and may your drams be as smooth as a badger's belly. #JohnnieWalkerBlack #WhiskyTales #SipResponsibly



### Conclusion

### Expectations

- Productivity Increase
- Redefine the work of leaders as much as employees
- GenAl will transform transformation
- Improve existing products and services
- Encourage innovation and growth
- Uncover new ideas and insights

### Concerns

- Estimating and demonstrating Al value
- Lack of talent/skills
- Lack of confidence in results
- Intellectual property issues
- Misuse of client or customer data
- Ability to comply with regulations
- Lack of explainability / transparency

### How In-DB LLMs address concerns

- Data does not leave database for embeddings
  - Client and customer data stays within corporate boundaries
- Oracle 23ai supports rapid prototyping of LLM-enabled features with existing infrastructure and skills
- The LLM will run on existing infrastructure and will not lead to higher cost in operations (things like backup and HA are already in place)
- Open Source models solve many issues
  - Transparency of model source and training data
  - No client data or internal data used for training of proprietary public models
  - Can be fine-tuned for specific use cases

# Thank you!

, Get the Slides



maximillian.jahn@orbitcloudsolutions.at max@maxjahn.at



https://blog.maxjahn.at



https://www.linkedin.com/in/maximillianjahn/



https://github.com/maxjahn

