


DBSQL AI FUNCTIONS



Jeremy Lewallen
Last updated June 2024

Product safe harbor statement

This information is provided to outline Databricks' general product direction and is for informational purposes only. Customers who purchase Databricks services should make their purchase decisions relying solely upon services, features, and functions that are currently available. Unreleased features or functionality described in forward-looking statements are subject to change at Databricks discretion and may not be delivered as planned or at all

The Next Opportunity: Analysts WANT to use Machine Learning

- **Demand:**
Analysts are eager to leverage ML for deeper insights, better analysis, and easier workflows.
- **The Problem:**
 - No easy, batch way to use ML across 1000s of rows.
 - At best, requires 3P tools or other languages (ie..Python)



Goal: Deliver access to AI that ANY analyst can use regardless of AI skill level

If they know SQL, then they should be able to use AI in DBSQL

- Ready-to-use models in SQL syntax
- Use SOTA models
- Infrastructure, tuning, and fitting all created by Databricks' ML Engineering and data scientists.
- AND....Fully managed and out-of-the-box



The solution is DBSQL AI Functions

Out-of-the-box suite of functions to call ML Models

THE DETAILS

- Range of out-of-the-box use cases
- Expanding quarterly
- Fully managed functions
- Used by Databricks' largest customers

Custom Models



Text Classification



Sentiment Analysis



Sensitive Data Masking



Text Similarity



Text Summarization



Language Translation



Information Extraction



Demo Use Case: How do users feel about AI Functions?

The screenshot shows a database interface with a query editor and a results table. The query is: `select * from main.default.ai_tweets_2024;` The results table has columns for ID, Date, and Tweet. The results show 7 rows of tweets related to DBSQL's AI functions.

ai forecast • New query Observability • New query • +

▶ Run all (1000) ▼ main.default Query (Preview): OFF

```
1 select *
2 from main.default.ai_tweets_2024;
3
4
5
6
7
8
9
10
11
12
13 /*
14
```

Raw results Line 1 Line 1 + New result

	ID	Date	Tweet
1	1	2024-05-27	DBSQL's vector search has transformed my research by locating essential papers in no time! #DBSQL #Academic
2	2	2024-04-21	DBSQL's sentiment analysis improves our understanding of customer feedback. #DBSQL #CustomerService
3	3	2024-05-13	Using DBSQL's ai_query to find top store products was straightforward! #DBSQL #Retail
4	4	2024-02-20	DBSQL's forecast function predicts a fantastic sales quarter ahead! #DBSQL #Business
5	5	2024-03-02	Using DBSQL's ai_query to locate top products in my store was a breeze! #DBSQL #Retail
6	6	2024-04-28	DBSQL's ai_query function made identifying top-selling products in my store easy! #DBSQL #Retail
7	7	2024-03-08	Our grasp of customer feedback has greatly improved thanks to DBSQL's sentiment analysis! #DBSQL #CustomerService

366 ms | 100 rows returned

ANNOUNCEMENTS



3 ANNOUNCEMENTS

Expanding Analysts' capabilities with faster and more intelligent functions



10x Speed Increase

- Calling custom models and LLMs are now 10x faster



Time Series Forecasting

- Built in time series forecasting that evaluates multiple models to find the best fit for your data



Vector Search

- Mosaic AI best-in-class vector search easily accessible via DBSQL

QUERYING CUSTOM MODELS & LLMS, BUT 10x+ FASTER

Unparalleled performance with our new multi-threading and batch API capabilities, delivering a 10x boost in query speed and efficiency

- Parallelization
- Batch API for handling multiple rows at once
- Multi threading support

```
SELECT ai_query(  
  'databricks-mixtral-8x7b-instruct',  
  CONCAT('Could you tell me the state name of the zip code provide  
) AS output  
FROM (  
  SELECT *  
  FROM samples.nyctaxi.trips  
  LIMIT 100  
)
```

ROLLING OUT NOW TO
GENERAL AVAILABILITY

Demo Use Case: What use cases are customers tweeting about?

The screenshot shows a SQL query editor interface. At the top, there are tabs for 'ai forecast', 'New query', 'Observability', and another 'New query'. Below the tabs, there are buttons for 'Run all (1000)', a dropdown menu, and a 'Query (Preview): OFF' indicator. The main area contains a SQL query:

```
1 select
2 ai_analyze_sentiment(Tweet) as sentiment,
3 ai_classify(Tweet, array('forecast', 'ai query', 'vector search', 'other', 'summarize', 'classify')) as function_used
4 from main.default.ai_tweets_2024;
```

Below the query editor, there are tabs for 'Raw results', 'Line 1', and 'Pivot table 1'. The 'Pivot table 1' tab is active, displaying a pivot table with the following data:

forecast	positive	19.0
other	neutral	7.0
	positive	11.0
summarize	positive	22.0
vector search	positive	20.0
Totals		100.0

7 rows

ANNOUNCING: Mosaic AI Vector Search in DBSQL

Mosaic's best-in-class data retrieval is now out-of-the-box. Find similar records and integrate with large language models seamlessly.

- Perform state-of-the-art KNN searches across your data index
- Vector Search for up to 100 million records
- Easy out-of-the-box RAG

```
SELECT *  
FROM VECTOR_SEARCH(  
    index => "main.db.my_index",  
    query => "iphone",  
    num_results => 2  
)
```

NOW ENABLED IN PRIVATE
PREVIEW

Demo Use Case: What do retail users think about AI Query?



ANNOUNCING: Built-in Time Series Forecasting

Multiple models evaluated on-the-fly, to extrapolate time series data, handling grouped, multivariate, and mixed-granularity data efficiently.

- Forecasting at scale with simultaneous evaluation of multiple models.
- Multiple built-in tuning parameters
- No pre model training required

NOW ENABLED IN PRIVATE
PREVIEW

```
WITH aggregated AS (  
  SELECT  
    DATE(pickup_datetime) AS datetime,  
    SUM(fare_amount) AS revenue  
  FROM  
    nyctaxi.trips  
  GROUP BY  
    1  
)  
  
SELECT *  
FROM AI_FORECAST(  
  TABLE(aggregated),  
  horizon => '2016-03-31',  
  time_col => datetime,  
  value_col => 'revenue'  
)
```

Demo Use Case: Forecast revenue for the next ~ 2 mo



GA AND PREVIEWS STARTING NOW



10X SPEED INCREASE
ROLLING OUT TO GA NOW!



VECTOR SEARCH IS IN
GATED PUBLIC PREVIEW



TIME SERIES
FORECASTING IS IN
PRIVATE PREVIEW



ROADMAP

Enabling analysts to deliver AI driven & high quality insights



From Fast to Blazing Fast

Access your custom models, out-of-the-box models, or LLMs faster than any other production OLAP



Simplified User Journeys

Functions to call ML models rather than special SQL syntax



Explaining model results

Understand why the result is being generated and what the model is doing

APPENDIX

