Revolutionizing Insight Delivery with AI

CDAO 2024

Wessel de Meyer – Senior Sales Engineer

Data Literacy

Skills Shortages

Business Demand

Technical Debt

Rate of Innovation

Improving Analyst Productivity with GenAI

Our Guiding Principles for GenAI

EFFICIENT

EASY

TRUSTED

Snowflake Copilot

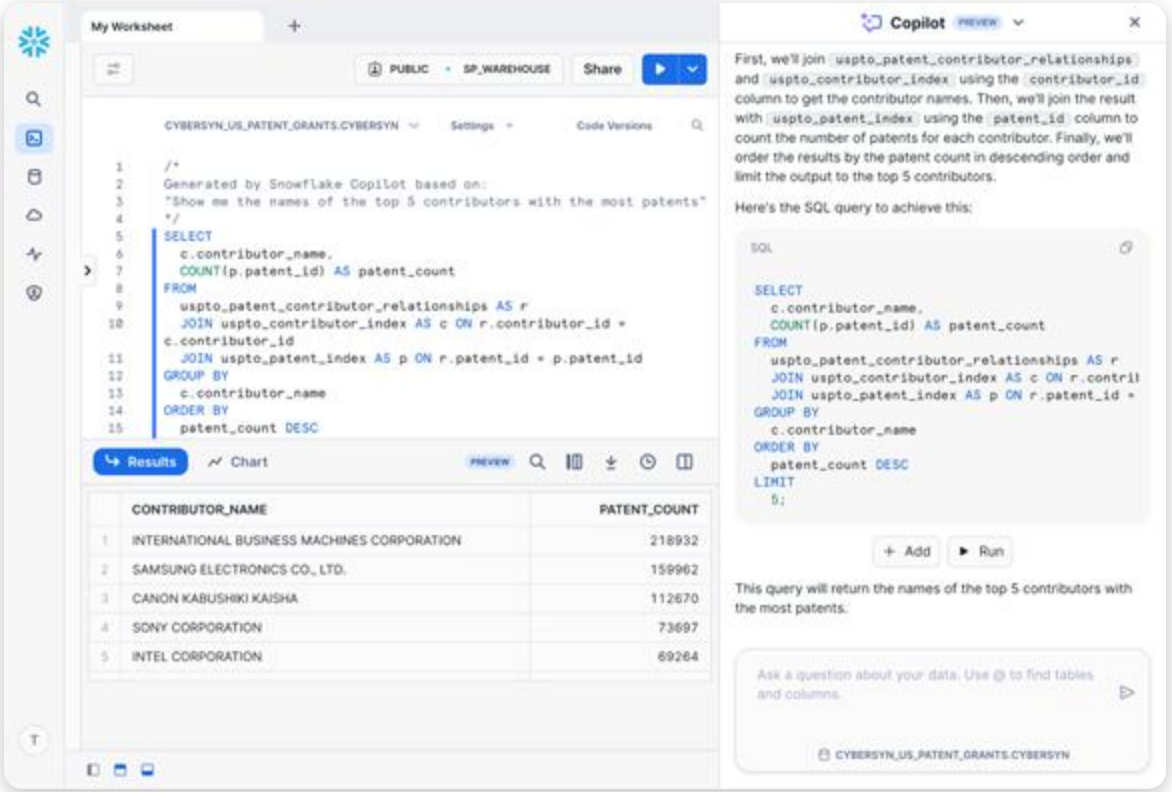

Snowflake Copilot helps **analyst** and **developers** to **explore datasets, generate SQL queries** and **answer questions** through Snowflake documentation.

Notebooks COPILOT

Packages Active Run all

Python as cell1

```python
# Import python packages
import streamlit as st
import pandas as pd

# We can also use Snowpark for our analyses!
from snowflake.snowpark.context import get_active_session
session = get_active_session()

```

SQL as cell2 0.7s

```sql
-- Welcome to Snowflake Notebooks!
-- Try out a SQL cell to generate some data.
SELECT 'FRIDAY' as SNOWDAY, 0.2 as CHANCE_OF_SNOW
UNION ALL
SELECT 'SATURDAY',0.5
UNION ALL
SELECT 'SUNDAY', 0.9;
```

	SNOWDAY	CHANCE_OF_SNOW
0	FRIDAY	0.2
1	SATURDAY	0.5
2	SUNDAY	0.9

cell1
cell2
cell3

Ask Copilot

WDM

Feedback

Today’s Process is Slow and Inefficient

Business Users

Know the business well, but writing SQL queries is not part of their workflow. Use BI dashboards that may be out of date or ask questions to data teams.

Data Questions

Ad Hoc Answers

Data Teams

Build and maintain BI dashboards, which requires significant effort. Receive ad hoc data questions from business users, which can be a time sink.

Cortex Analyst: Self-Service Data Analytics

Built With Meta's Llama 3 and Mistral Large Models

Business Users

Can answer their data questions via a self-service workflow using a conversational interface. Able to move faster and make data-driven decisions.

Data Teams

Set up the semantic model and let Cortex Analyst take care of the rest. No longer receive inbound requests, which helps save valuable time.

How Does Cortex Analyst Work?

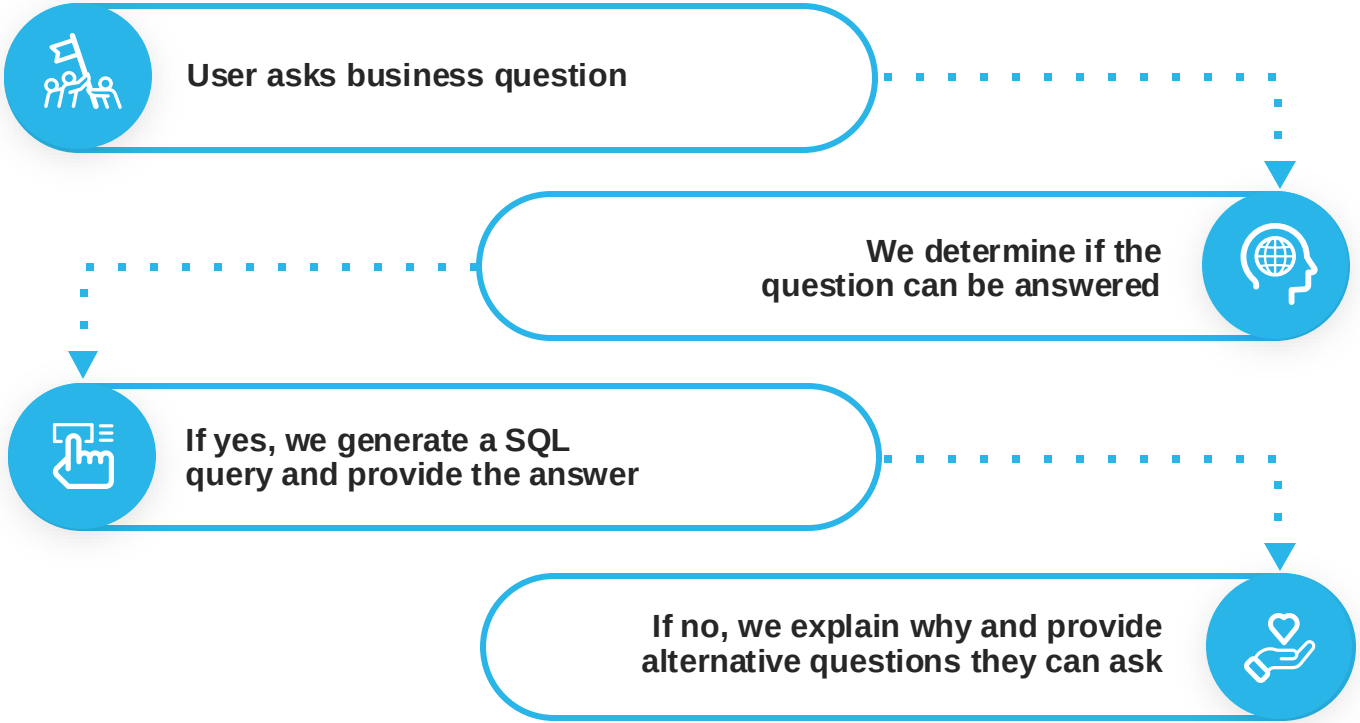










Streamlit Apps CORTEX_ANALYST

Active Share Edit

Cortex analyst

Semantic Model: `revenue_timeseries.yaml`

What is your question?

Cortex Analyst

Business Benefits

- Shorten the time required to deliver answers to the business
- Reduce the demand on report analysts by minimizing the need for them to create reports and charts for every information request
- Simplify end user experience by embedding Cortex Analyst in front-end applications

Intelligent Document Processing with GenAI

Universal Across Industries

Public Sector

Citizen Service

Application Processing

Healthcare

Patient Care

Claims Processing

Financial

Loan Processing

Tax Compliance

Retail / CPG

Supply Chain Optimization

Product Management

Manufacturing

Quality Assurance

Inventory Management

Technology

Contract Management

Product Development

Advertising / Media

Contract Negotiation

Content Licensing

Pitfalls of Manual Document Processing

Error Prone

Manual processing frequently results in errors due to human oversight.

Tedious and Slow

Manual processing is repetitive and an inefficient use of valuable time.

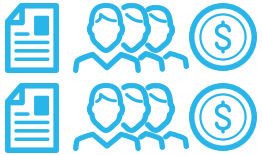

Resource Intensive

Manual processing required substantial manpower and investment.

Harnessing Intelligent Document Processing (IDP)

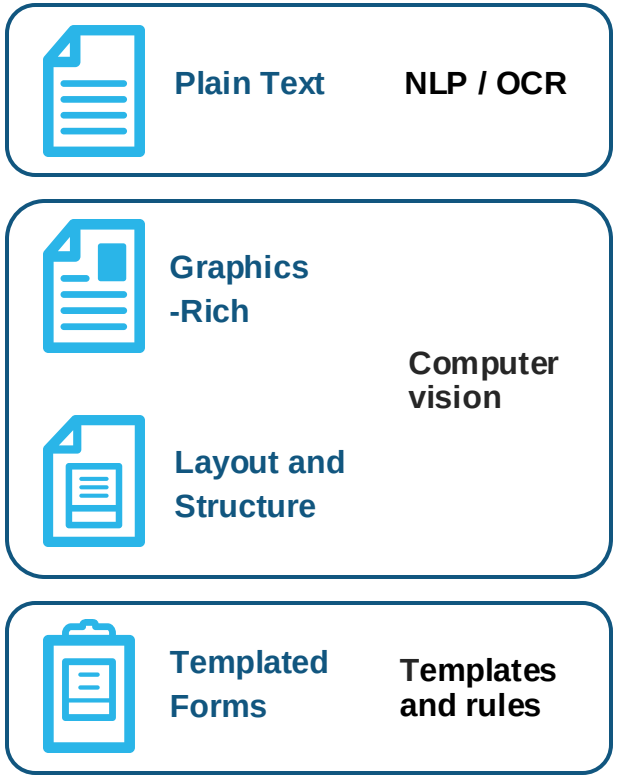

Intelligent Document Processing

IDP streamlines the entire document lifecycle. It leverages tools like OCR, NLP and computer vision to efficiently manage and process unstructured data.

Harnessing Intelligent Document Processing (IDP)

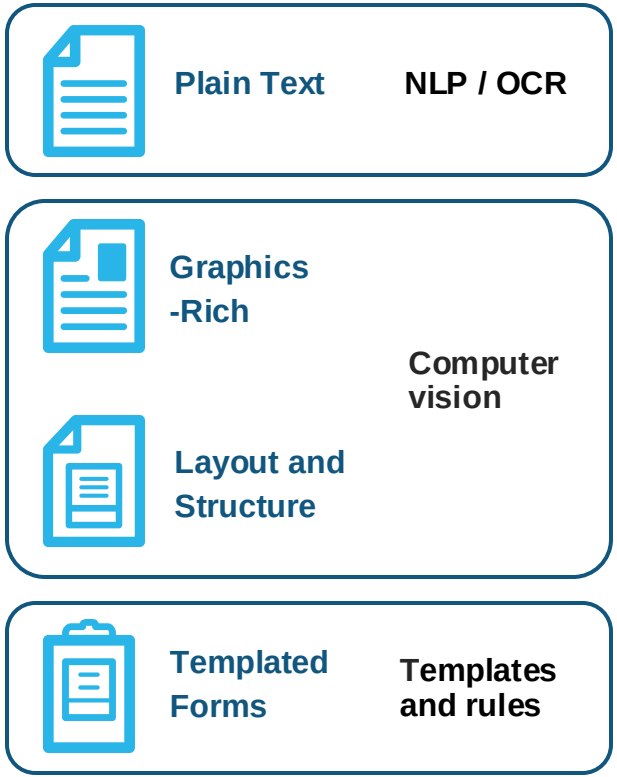

Intelligent Document Processing

+ GenAI

IDP streamlines the entire document lifecycle. It leverages tools like OCR, NLP, computer vision, AI and LLMs to efficiently manage and process unstructured data.

INTRODUCING DOCUMENT AI

SNOWFLAKE WORLD TOUR

Document AI

WHAT IS IT

New workflow utilizing a Snowflake LLM for efficient extraction of content from unstructured documents

WHY USE IT

Optimize your business with higher efficiency, lower cost of manual labor and lower human error, with an industry-leading LLM

HOW TO USE IT

In a natural language interface, ask questions to test model, fine-tune with a single-click if needed, and execute field extraction across multiple documents via SQL. Store outputs in Snowflake tables for downstream consumption.

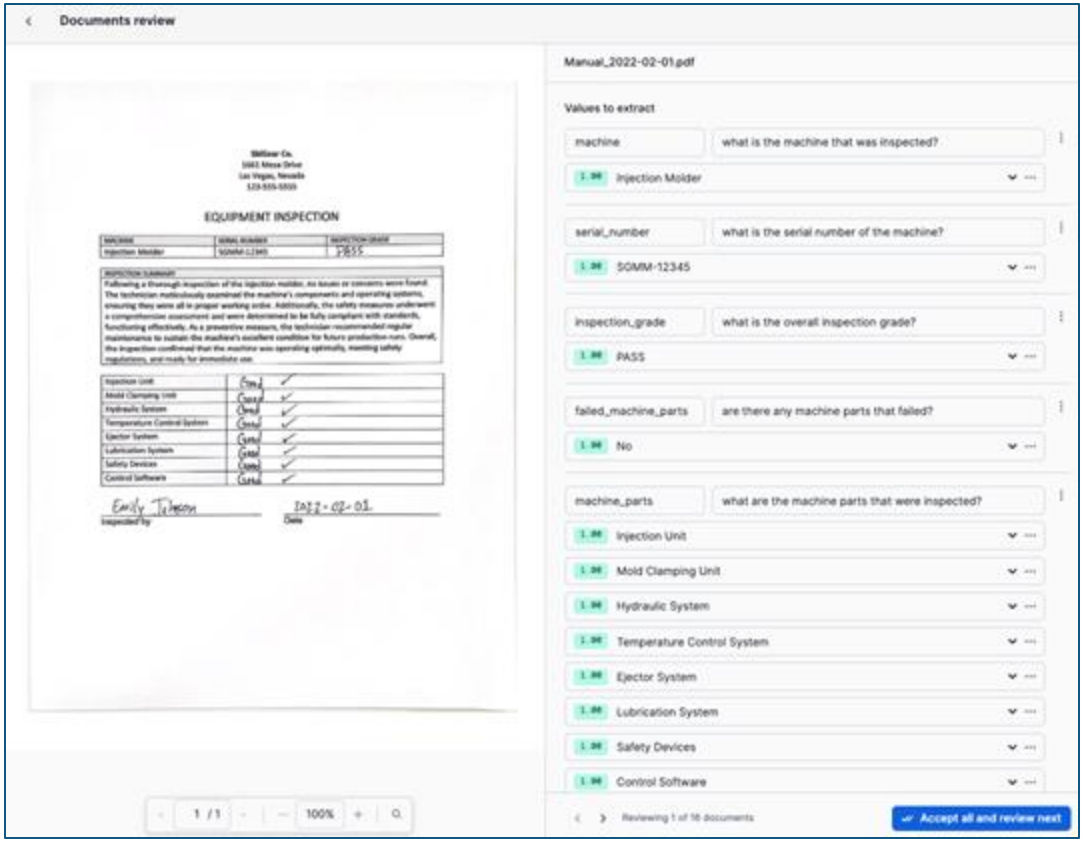

Document AI Document Processing

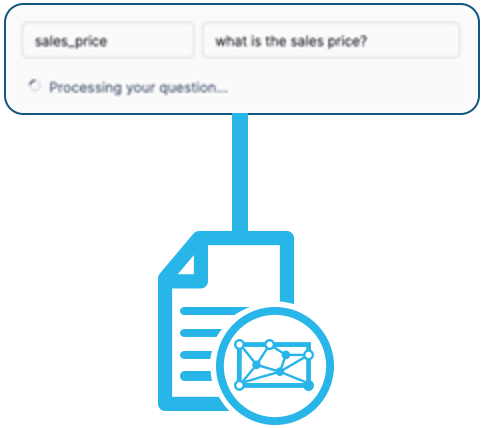

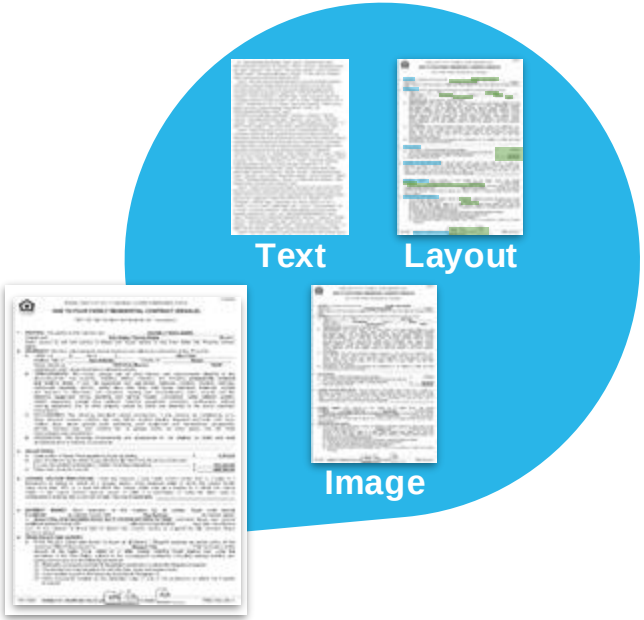

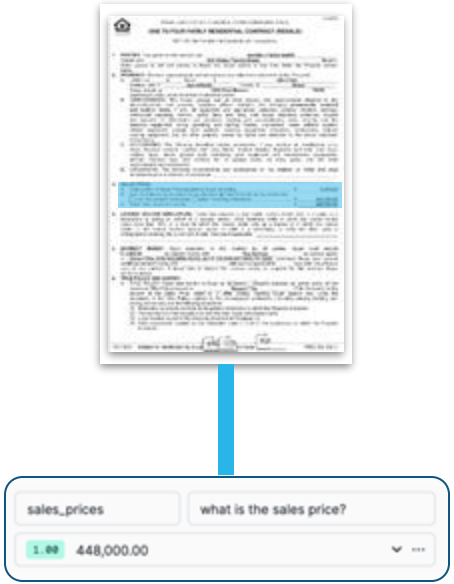

Prompt Interpretation

Document Parsing and Inference

Response Generation

Snowflake Arctic-TILT

Arctic-TILT

Text Image Layout Transformer

A purpose-built, multimodal large-language model of **0.8B parameters** that fits in a single A10 GPU machine

Benchmark

90.20 **ANLS**

DOCVQA

A standard in Visual Question Answering Benchmarking

Achieves top results in benchmark standards requiring layout understanding.

GPT-4 = 88.4

User Experience

Document Owner and **Stakeholder**

Natural Language User Interface

Define

Train

Evaluate

Publish

SQL

Load Documents

Automate

Run Inference

Data Engineer

Consumption

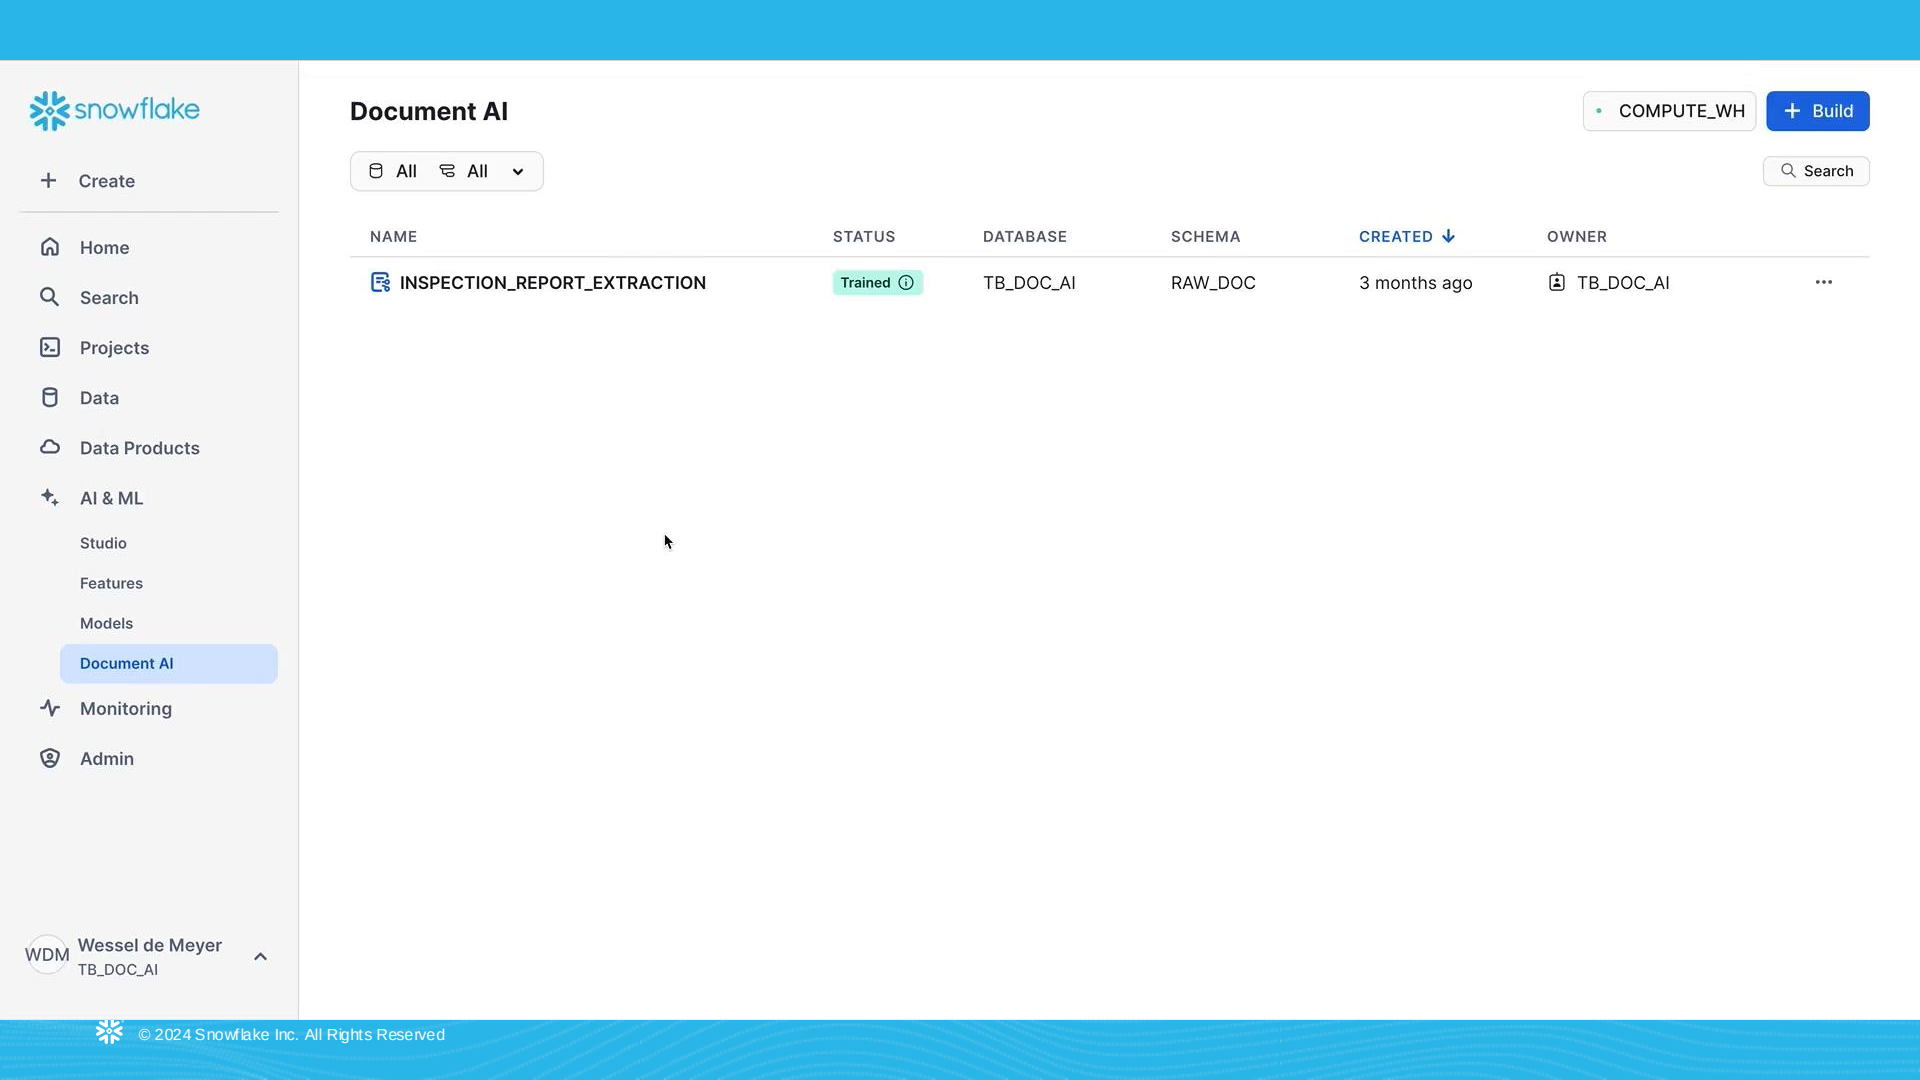
snowflake
Create
Home
Search
Projects
Data
Data Products
AI & ML
Studio
Features
Models
Document AI
Monitoring
Admin
WDM
Wessel de Meyer
TB_DOC_AI
Document AI
COMPUTE_WH
Build
All
All
Search
NAME
STATUS
DATABASE
SCHEMA
CREATED
OWNER
INSPECTION_REPORT_EXTRACTION
Trained
TB_DOC_AI
RAW_DOC
3 months ago
TB_DOC_AI
© 2024 Snowflake Inc. All Rights Reserved

Health Partnerships – POC

- **14 000 Procurement contracts across the health system**
- **Document AI POC Scope**
  - A total of 19 contracts of average 25 pages in PDF format
  - 15 PDF documents were used to train the model based
  - 4 PDF documents were reserved for validation

Value & Questions

	VALUE NAME	QUESTION
1.	**effective_date**	What is the effective start date of this contract in the format dd-mm-yyyy?
2.	**expiration_date**	What is the expiration date of this contract in the format dd-mm-yyyy?
3.	**amount**	What is the total amount in this contract?
4.	**currency**	What is the currency used in this contract?
5.	**parties**	Who are the parties involved in this contract?
6.	**contract_description**	What is the description of this contract?
7.	**contract_ref_number**	What is the contract reference number of this contract?
8.	**is_variation**	Is this a variation of a contract?
9.	**variation_number**	What is the variation number of this contract?

Results Summary: 86.11% Accuracy

Te Whatu Ora
Health New Zealand

Why Document AI

Proximity to Data

Minimize latency by processing where your data already resides

No ML expertise or APP Dev needed

Use a natural language interface with the ability to fine-tune models.

Purpose-built proprietary model

Utilize a cutting-edge proprietary model purpose-built for business document extractions

Fully Managed and Governed

Serverless. No need to deploy models, test them, secure them, upgrade them, or scale GPU clusters

Thank You