Snow Leopard

Bridging the Gap between AI and Business Data

Deepti Srivastava
Founder & CEO
snowleopard.ai

About me

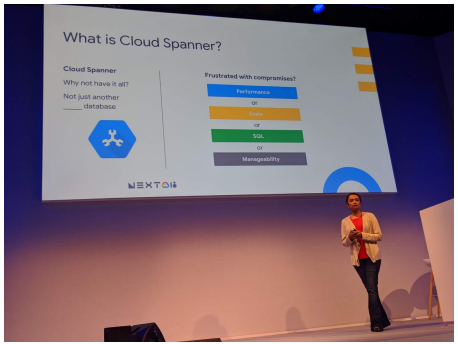

~2 decades in data platforms and infrastructure
part of 2 platform shifts

Head of Product, Observable
data visualization & analytics, BI

Founding PM, Google Spanner
cloud and data infrastructure-as-a-service

Engineering Manager, Oracle RAC DB kernel
technical infrastructure, distributed systems

M.S. in Computer Engineering, Carnegie Mellon
distributed fault tolerant systems

So what's the problem?

AI is changing the world and changing the way people work

BUT …

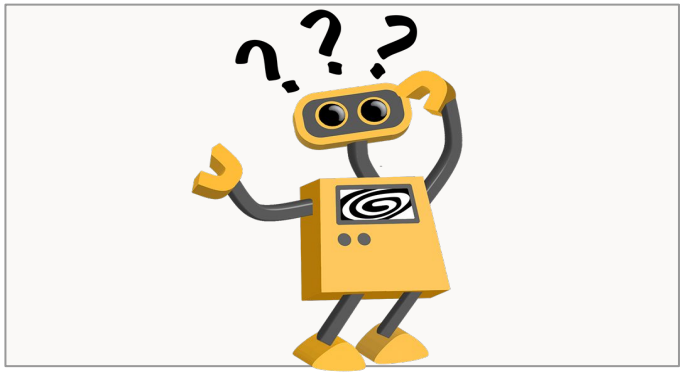

Are we able to solve real-world problems and automate complex decision-making?

“I just want to know my order status!”

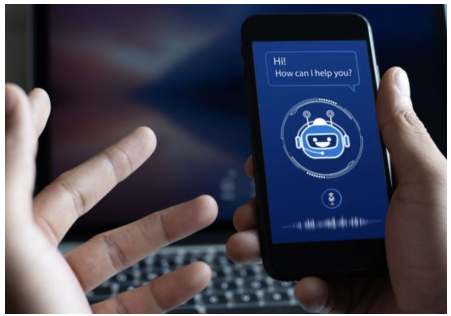

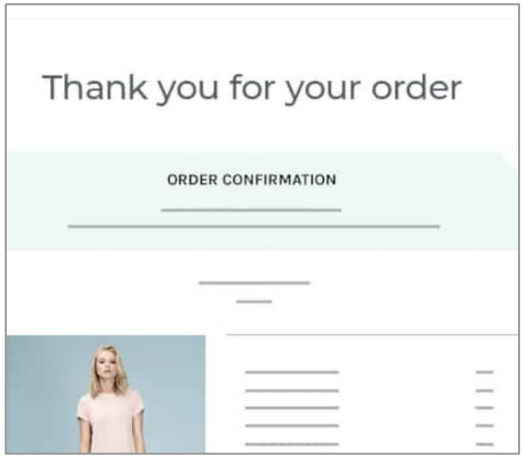

Other *live-data* use cases that require human-in-the-loop

“Did I order Pad Thai in the last week?”

Are more umbrellas on their way now because its raining and we sold out unexpectedly?”

What's missing?

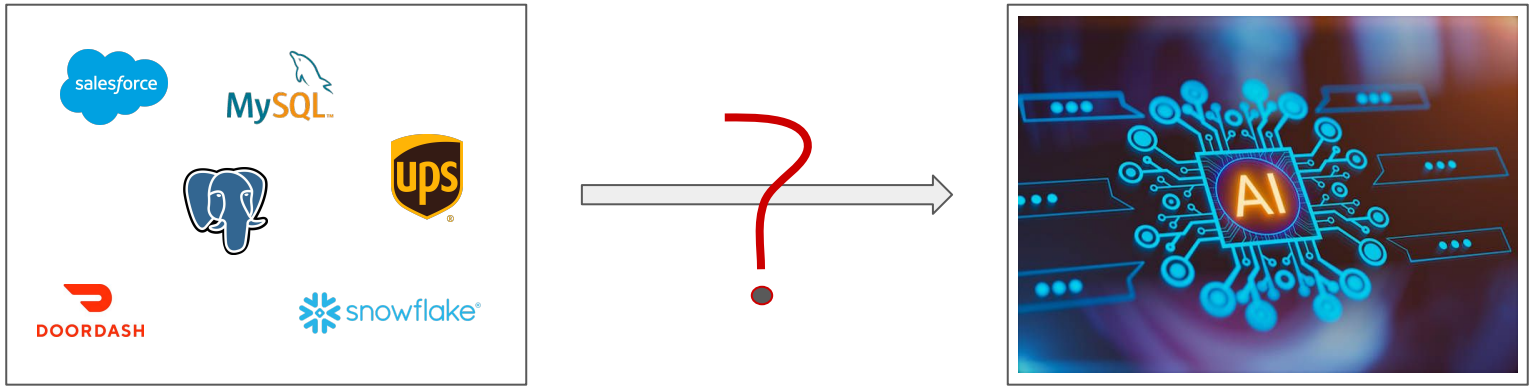

Access to the right data from the right source at the right time

Current solutions

What are the current best-practice solutions in the market to bridge the gap between business data and AI

Fine-Tuning LLMs

COMMON USE CASES

Domain-specific workloads: Understanding medical/financial terminology

Improving qualitative aspects: Font, tone, verbiage for external responses

CONSIDERATIONS

Needs **accurate**, **targeted**, **well-labeled datasets** (lots of data)

Sometimes prohibitively **expensive**

Requires expertise to prep data and do the fine-tuning

BEST FOR

Static or mostly static **data** where **data freshness** and **liveness aren't critical** to decision making

RAG

COMMON USE CASES

Connecting knowledge bases to LLMs for enhanced Q&A systems

Automating retrieval workflows of (multi-modal) **unstructured data**

Summarization, **classification** and related conversational tasks

CONSIDERATIONS

Staleness, loss of freshness: inherent in ETL & data pipelines

Loss of accuracy: extracting from DB/APIs to a fuzzy matching system

Expensive: data in motion, complex systems, hard to maintain

BEST FOR

When **data freshness and accuracy aren't critical** to the workflow or for decision making

AI Agents

COMMON USE CASES

Specialized conversational AI agents: personal assistants like Siri

Autonomous robots in specific domains managing a given set of tasks

CONSIDERATIONS

Purpose-built: specific goals/rules

Can be **extremely complex** to build
Requires **specialized ML expertise**

Expensive to build and train and **not scalable**

Still emerging field: lots of unknowns

BEST FOR

Automating **specific**, **predefined tasks** and **workflows** when you **have data+compute+expertise** to setup and manage

But….there is STILL a gap

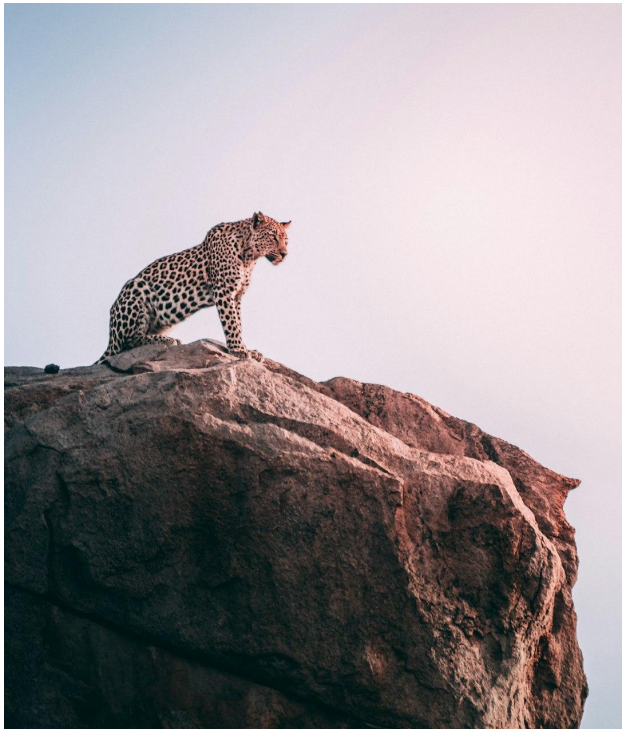

There is still nothing that makes **live business data** available to the AI-powered systems **directly**

…it's only done in very **narrow ways** that are **complex** and **unscalable**

Snow Leopard

Bridging the gap between live business data and AI systems to realize the full potential of AI for your business

Solving the “live data” problem

- Retrieve live data from any data source on-demand at query time
- Do away with repeated ETLs and complex data pipelines
- Provide simple API for AI developers
- Make management simple

AIM: Simplify AI workflows and unlock business use cases

Why is it different

We are building a data retrieval system that has **intelligence about the data** instead of predefined pipelines and workflows

Query data source natively and in-situ instead of unnecessary ETL and managing complexity of data in motion

Benefits

- Accuracy → get freshest data needed at query time
- Better performance → query source directly and natively
- Cheaper → do away with unnecessary, complex ETL and data pipelines
- Opens up more use cases → live data from any data source

We are early but if you want to learn more come talk to us!

deepti@snowleopard.ai

Thank You!

hello@snowleopard.ai