



Think Deeper with Claude

Building smarter Copilots with Copilot Studio and MS Foundry

Mitul Tailor & Sarang Brahme | AICO NYC | June 26th, 2026

WITH GRATITUDE

Thank you to our sponsors

PLATINUM



Microsoft
Solutions Partner

GOLD

SOHO
DRAGON

SILVER

withum+

COMMUNITY

AI Maitri



**Built by the community,
backed by our partners**

A volunteer-run conference is only possible
because these organizations believe in it.
Thank you for making it happen.

Sponsor a future event · aicoevents.org/sponsors

About Me

Mitul Tailor

Sr. Data Engineer & Consultant

Long View Systems – Toronto, Canada

Microsoft Certified Trainer (MCT)

Delivering official Microsoft training in data & AI

Passionate Technologist

Data, generative AI, and enterprise Copilots

Agentic AI Builder

Autonomous agents with Copilot Studio & Microsoft Foundry

Community Contributor

Sharing on Claude, Copilots, and applied AI

Let's Connect


LinkedIn: [@mitultailor15](https://www.linkedin.com/in/mitultailor15)



About Me

Sarang Brahme

[linkedin.com/in/sbrahme](https://www.linkedin.com/in/sbrahme)



Almost 2 Decades in the Tech Trenches

A Principal AI Architect, Technical Product Owner/Leader who turns big ideas into real, working AI roadmaps.

C-Suite and VC Translator

I speak fluent "Business Value," "Data Science," and "Software Engineer"—and help them all get along.

Current Obsession

Agentic AI, MLOps, building/advising startups. If it makes life easier intelligently, I'm building it.

Tech Stack of Choice

Azure at the moment, but also - Databricks, Fabric, GCP, Open Source

Speaker & Practitioner

I do talk about AI on stages; and I actually build the tech when I get off stage.

Warning

Once started, I can't stop talking about agentic AI, quantum computing, heavy metal, and philosophy.

Agenda

01



02



03



04

The
Background

Foundations &
Design

Intelligence &
Agents

Claude &
Copilot Studio

Feedback to Action

Engineer this!

NLP Core

Meet Claude

What are Verbatims?

Reference Architecture

Objective Mapping

Copilot Studio

What do clients 'need'?

Agentic AI

DEMO!!

01

Background

From Guest Feedback to Business Action

“Every guest leaves a signal. Most brands only hear the noise.”



Turn every review into insight, every insight into action, and every action into a better guest experience.



What we mean by 'verbatim' reviews

verbatim

🔊 [ver-**bey**-tim] / ver'bei tɪm /

The raw, unedited free-text a customer actually wrote (app review, survey open-end, support transcript, social post)

Venue Rating



example # 1

Came here because of the high rating. It was busy, so the **delay in greeting** was understandable. We ordered two omelettes, but the wait was long—about 20–25 minutes for the first, and another 4–5 minutes for the second.

Food tasted good, but the plates looked quite worn and marked. Overall, decent food but service timing and presentation could be better.

example # 2

If you look at the menu, the theme of the 90s is there, but in **decor or in the music, you can not feel it.** My experience started that I got an "outsider" table. It is a table where basically coming a border in service between 2 waitress and you hidden next to the wall. So, one person is **not paying attention** to you because it is not her table, and your waitress simply can not see you. I'm not much picky about it. Hence, **I was okay with my book.**

Food

Service

Ambience

Overall

Parking

.....

What do Clients actually need?

01. Deep Decomposition

02. Strategic Mapping

03. Structured Persistence

04. Contextual Enrichment

05. Orchestration & Governance

06. Reframe!

Reframe:

Old framing: How positive is this text?



New framing: Which 'objective' is affected, in which direction, with what confidence, and what should happen next?

Output we want isn't a score — it's a routed, contextualized, auditable signal

02

Foundation & Design

Engineer this!

Data Engineering

- Pipelines for daily data load
- Landing, Cleaning, Aggregation
- Logging and monitoring

Agentic AI Framework

- Multi-AI agents:
- Semantic Kernel and Agent Framework

Natural Language & Text Analytics


- Reviews and verbatim surveys
- NLP ML models – mapping text into sentiments, objectives and scores

UI/UX

- Custom App or Copilot Studio
- Claude for:
 - Processing
 - Routing
 - Orchestration
- UX Pathways

DevOps and IaC

- DevOps pipelines – env control
- Integration and deployment pipelines




Reference Architecture on Azure



03

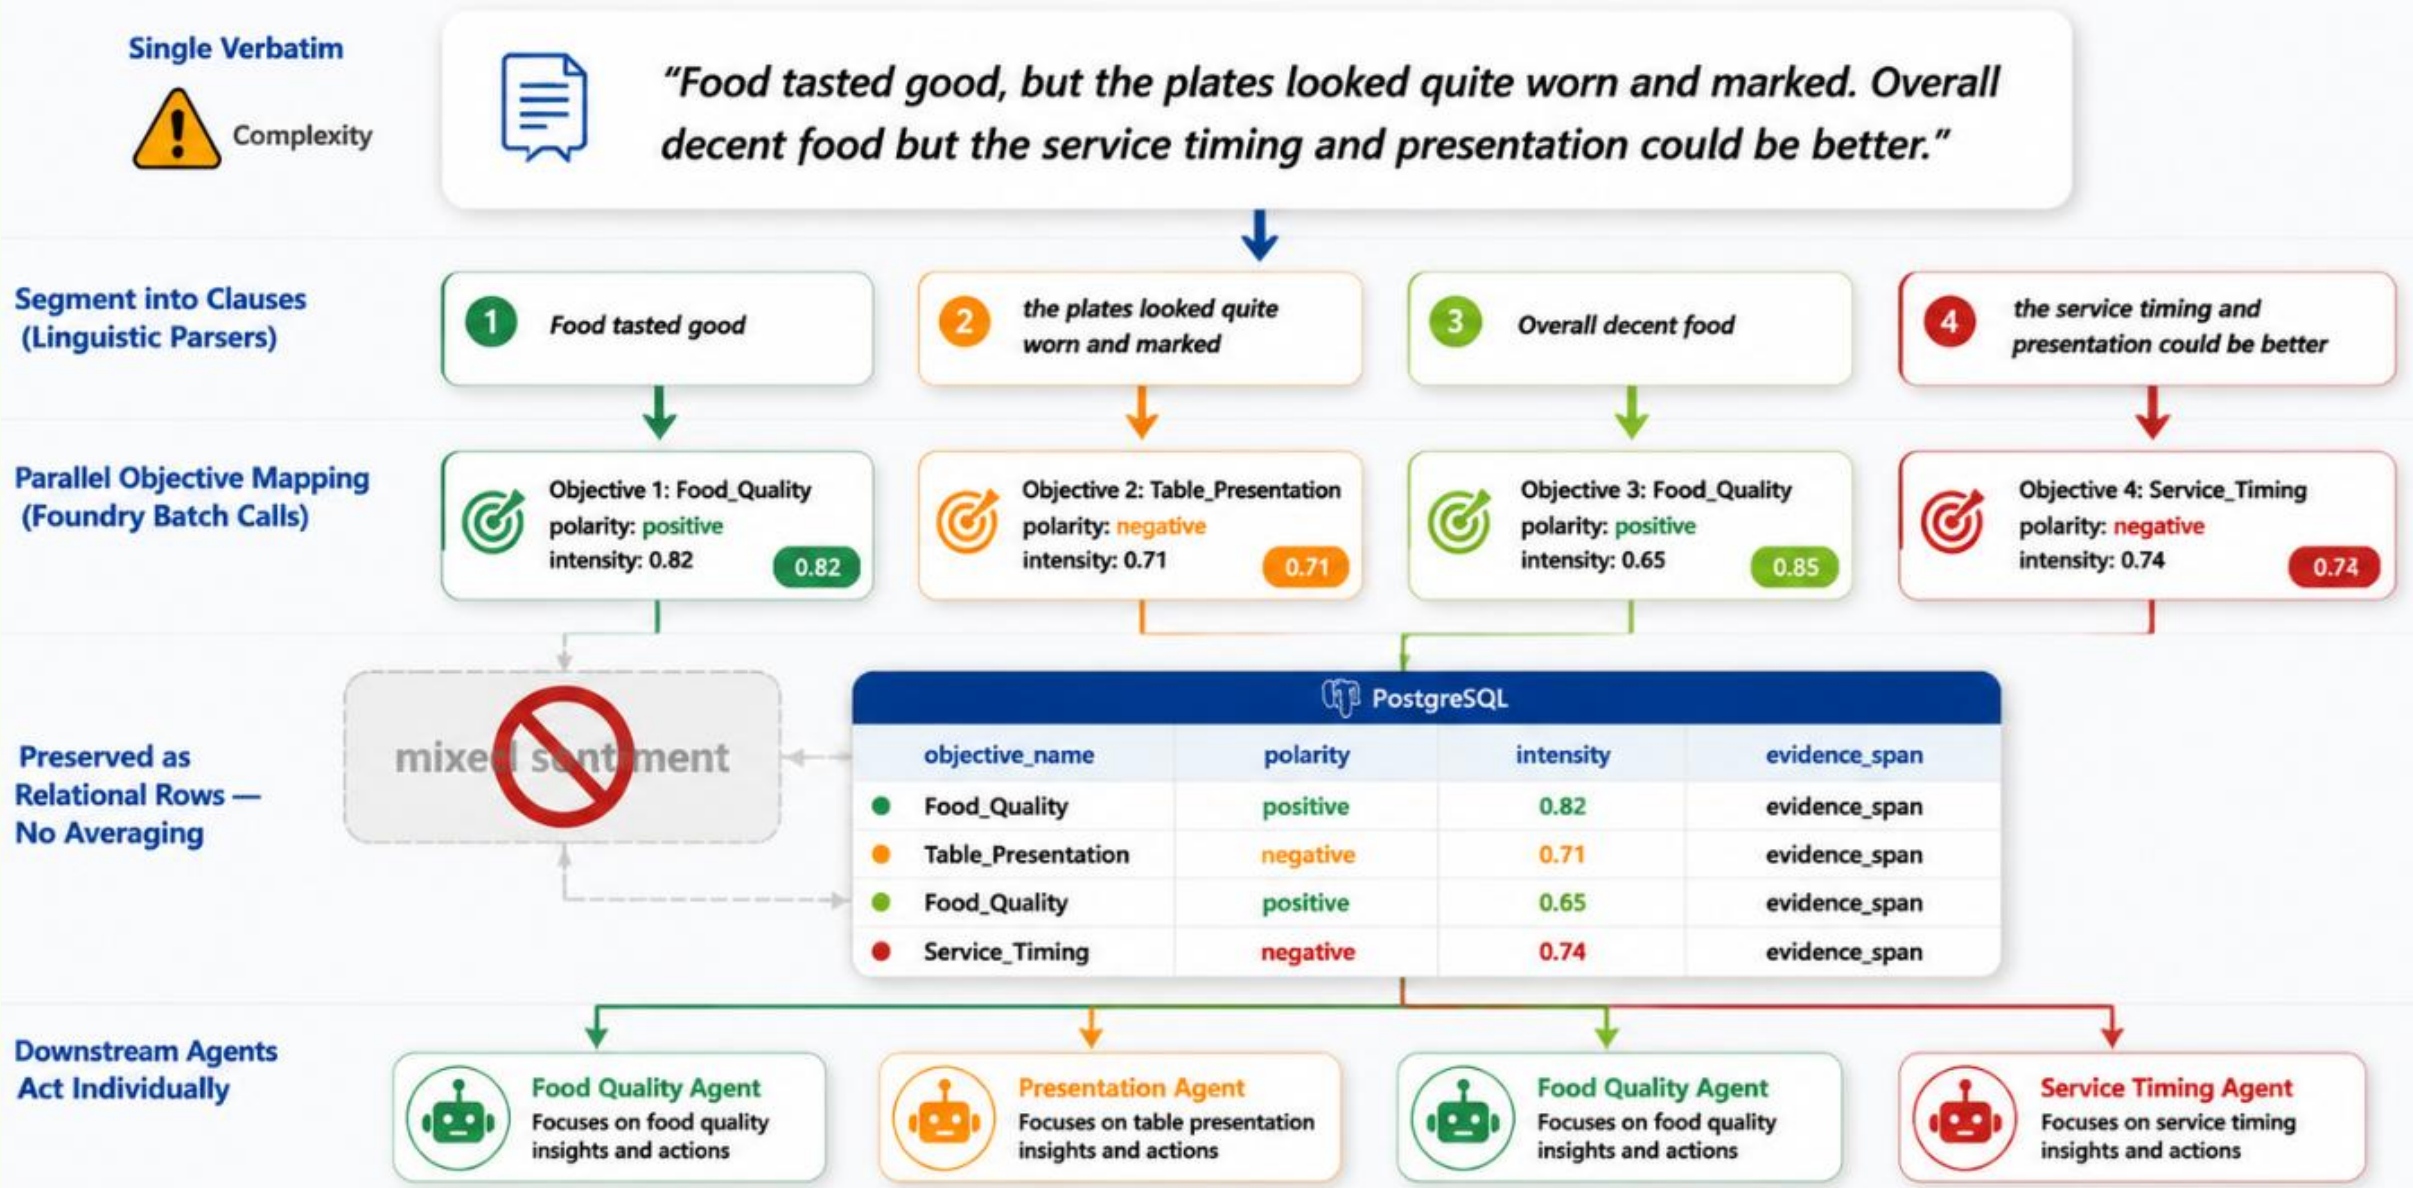
Intelligence & Agents

Semantic Objective Mapping - NLP Core




- Objective taxonomy is maintained as a master table in Postgres
 - Each objective has a human-readable definition + few-shot example clauses
- Mapping logic (Azure AI Foundry-hosted):
 - Retrieve top-k similar objective templates via vector search
 - Feed clause + retrieved templates into a GPT-4o / GPT-5 reasoning call
- Structured output: JSON with objective, polarity, intensity, rationale

Multi-Sentence Complexity



Why Agents? – From Analysis to Action




04

Tools & Demo

Anatomy of an AI agent

Every agent loops over the same core parts: a reasoning model guided by instructions, calling tools and memory until the task is done.



Retrieval isn't reasoning

Basic Q&A bot

Shallow, keyword-driven

- Returns shallow, fact-retrieval answers
- Short context windows miss meaning in long documents
- Nuanced or multi-part questions yield incomplete replies
- Cannot correlate data across sections of a report

Reasoning-first Copilot

Thinks, not just finds

- Sustains context across long documents and many turns
- Performs deep analysis of complex, cross-sectional data
- Answers multi-step, nuanced business questions reliably
- Surfaces insight, not just matching keywords

What Agents Look Like in a Demo

Copilot Studio

Dataverse/CSV

What Agents Look Like in Production

Copilot Studio

Orchestrator

Hotel Agent

Logistics
Agent

Quotation
Agent

Static Web App / API

DirectLine
SDK

Speech

Foundry
Summarize

Dataverse
Proxy

Microsoft Foundry

Destination Expert
(gpt-5.4)

AI Search

Code Interpreter

Power Automate Flows

Search Hotels

Lookup Reservation

Create Draft Quote

Data & Storage

Dataverse

Blob Storage

AI Search

External APIs

Sabre GDS

Expedia OTA

Google Maps
Platform

From script bots to AI agents

An agent uses an LLM to interpret intent, decide what to do, and act through tools — a step beyond rule-based chatbots.

Script-based chatbots

Fixed and rule-bound

- Driven by predefined scripts and rule-based logic
- Keyword / intent matching triggers fixed responses
- Handle guided, single-path conversations
- Require manual updates to maintain dialogue flow

Modern AI agents

Adaptive and reasoning

- Powered by LLMs such as Claude, GPT, and Gemini
- Natural, multi-turn conversations that hold context
- Handle dynamic tasks: summarize, reason, code, translate
- Trigger workflows and integrate with enterprise systems

Meet Claude – Anthropic's model

A frontier LLM built for long documents, complex reasoning, and nuanced language - now an agent engine inside the Microsoft ecosystem.

01

Claude Sonnet 4.6

Balanced daily driver

Best mix of speed and intelligence; the versatile default for coding, analysis, computer use, and vision, with a 1M-token context window.

02

Claude Haiku 4.5

Fastest, near-frontier

The fastest model with near-frontier intelligence-ideal for high-volume, latency- and cost-sensitive tasks.

03

Claude Opus 4.8

Deepest reasoning

Anthropic's most capable Opus-tier model for complex reasoning and long-horizon agentic coding, with adaptive thinking and a 1M-token context.

Copilot Studio goes multi-model

Makers can choose Claude as their agent's engine — no rebuilding, no new framework. Three steps in the familiar Copilot Studio experience.

01

Enable

Admin turns it on

Admins switch Claude on via the M365 Admin Center; on by default globally from Jan 6, 2026 (EU/UK/EFTA opt in).

02

Select

Pick per agent

Makers pick Claude in the model selector beside the default model — per agent or per skill.

03


Deploy

Ship with fallback

Agents run on Claude across Copilot Studio and M365, falling back to GPT-4o if it is disabled.

How a request reaches Claude

From the maker's view it's all M365 — Copilot Studio orchestrates the conversation and routes calls to Claude (or GPT) via Foundry.



Per-agent / per-skill model choice - Claude's response surfaces back inside Copilot; users never leave M365.

Where Claude shines for Copilots

Four capabilities make Claude a strong fit for knowledge-intensive, reasoning-heavy agent scenarios.

01 Long-form content mastery

Ingests and summarizes hundreds of pages in one pass — reports, contracts, knowledge bases.

02 Deep, multi-step reasoning

Breaks problems into steps and produces well-reasoned, structured answers.

03 Nuanced understanding


Interprets ambiguous or layered questions and responds with contextual relevance.

04 Multi-turn reliability

Retains earlier context so follow-up questions build a real analytical dialogue.

Scenario: the analyst Copilot

A Claude-powered Copilot acts less like a search engine and more like a human analyst — reading, reasoning, and recommending.



Business impact: automates deep analysis · unlocks strategic Q&A · retains context for true multi-turn investigation.

Demo Time!



Key takeaways

What to remember about bringing Claude into Copilot Studio.

01

Not just a model swap

Claude in Copilot Studio is a step-change in what your agents can understand and deliver.

02

Use it where it counts

Document analysis, long-form summarization, and multi-step problem solving are its sweet spots.

03

Choose per task

Multi-model means picking Claude or GPT per agent or skill — balanced, not absolute.

04

Design decides value

Thoughtful prompts and well-structured knowledge sources matter more than flipping a selector.

Next steps & questions

The capability is available today.

01

Get started

Enable & pilot

Enable Claude via Foundry and pilot one high-value agent.

02

Learn more

Explore the family

Review the model family and the Copilot Studio model selector.

03

Discuss

Plan a workshop

Bring your use cases — plan a hands-on workshop or pilot.

Thank you

AI Community Conference - NYC

2026 Session Evaluation

