



Figure 1: AI-Assisted Power BI for Business Analytics

Lab 1: Your First AI Conversation with Data

Module 1: The AI-Powered Analytics Landscape

Duration: 35 minutes **Dataset:** CloudRevenue (Microsoft Cloud Subscription Analytics) **Deliverable:** Reflection document (1 page)

Learning Objectives

By the end of this lab, you will be able to:

- Navigate Power BI Service and locate Copilot features
- Ask natural language questions about business data
- Apply the **CSAR Loop** for effective AI conversations
- Interpret AI-generated summaries and insights
- Recognize when Copilot excels vs. when it struggles
- Practice **appropriate reliance** — knowing when to trust vs. verify

Dialog Engineering: From Prompts to Partnerships

Most people approach AI like a search engine: type a question, get an answer, done. This is **prompt engineering** — optimizing a single input to get a better output.

But the most effective AI users do something different. They engage in **dialog engineering**: a structured, iterative conversation where both human and AI build understanding together.

“Dialog engineering replaces the monologue with a conversation. The user and the model build understanding iteratively, through questions, corrections, clarifications, and checkpoints.” — Correa, F. (2026). *Dialog Engineering: The Cognitive Science of Human-AI Partnership*

The CSAR Loop

The core protocol of dialog engineering is a four-phase cycle:

%%{init: {'theme':'base', 'themeVariables': {'primaryColor': '#e8f4fd', 'secondaryColor': '#e8f6ea', 'tertiaryColor': '#e8f6ea'}}}%%
flowchart LR

```

C[" CLARIFY<br/><i>Ask before acting</i>"]
S[" SUMMARIZE<br/><i>Verify understanding</i>"]
A[" ACT<br/><i>Execute the task</i>"]
R[" REFLECT<br/><i>Evaluate & learn</i>"]

```

C --> S --> A --> R --> C

```

style C fill:#f0f7ff,stroke:#c7d8ed,color:#3b82a0
style S fill:#fefcf0,stroke:#e8dbb8,color:#8b7830
style A fill:#edf6ee,stroke:#c2dbc8,color:#3d7a4a
style R fill:#f5f0ff,stroke:#d4c8e8,color:#6b4e9e

```

Phase	What You Do	Example with Copilot
Clarify	Ask questions to understand the task fully	“What time period does this data cover?”
Summarize	Verify shared understanding before proceeding	“So the report shows Q1-Q4 2024 revenue by region...”
Act	Execute the agreed work	Ask Copilot: “Show me revenue by region for Q4 2024”
Reflect	Evaluate the outcome and learn	“That answered the question, but I should verify the APAC number”

Why CSAR matters: A single prompt is a monologue. CSAR turns your AI interaction into a conversation where you iterate toward the right answer rather than hoping to get it on the first try.

Appropriate Reliance: When to Trust, When to Verify

Not every AI output deserves the same level of scrutiny. Use these three questions:

Question	If YES...	If NO...
Is the cost of an error high?	Verify carefully	Trust more freely
Does Copilot have a track record for this task?	Trust based on history	Verify until you build confidence
Can I verify quickly?	Always verify	Spot-check based on risk

Practical Rule: High stakes + No track record + Slow verification = Check everything. Low stakes + Proven track record + Fast verification = Trust the output.

The Copilot Conversation Model

Now that you understand dialog engineering, see how AI-assisted analytics differs from traditional BI:

```
%%{init: {'theme':'base', 'themeVariables': {'primaryColor': '#e8f4fd', 'secondaryColor': '#e8f6ea', 'tertiaryColor': '#e8f6ea'}}}
flowchart LR
```

```
    subgraph Traditional["Traditional BI"]
        A[ Build Report ] --> B[ Scan Visuals ]
        B --> C[ Calculate ]
        C --> D[ Document ]
    end

    subgraph AIAssisted["AI-Assisted Analytics (CSAR)"]
        E[" Clarify<br/>What do I need?"] --> F[" Summarize<br/>Ask Copilot"]
        F --> G[" Act<br/>Review Response"]
        G --> H[" Reflect<br/>Verify & Refine"]
        H -->|"Iterate"| E
        H -->|"Done"| I[ Insight Found ]
    end

    Traditional -. ->|"Evolution"| AIAssisted

    style Traditional fill:#f8f9fa,stroke:#d0d7de
    style AIAssisted fill:#f0f7ff,stroke:#c7d8ed
```

The key shift: from building reports to having structured conversations with data.

Before You Begin

Required Reading

1. **CloudRevenue BRD** - View BRD
 - Understand the business context: VP of Cloud Sales Operations
 - Review the 5 key business questions to answer
2. **CloudRevenue Data Dictionary** - View Data Dictionary
 - Familiarize yourself with available tables and columns
 - Note the relationships between Products, Regions, and Subscriptions

What You'll Need

- Web browser (Edge, Chrome, or Firefox)
- Power BI Service account
- Access to course workspace (provided by instructor)

Tip: Bookmark the Power BI Service URL (app.powerbi.com) — you'll use it every lab.

Lab Instructions

Part 1: Access the Report (5 minutes)

1. **Navigate to Power BI Service**
 - Open your browser and go to app.powerbi.com
 - Sign in with your university credentials
2. **Find the Course Workspace**
 - In the left navigation, click **Workspaces**
 - Look for the workspace named **VT-MBA-PowerBI-[Section]**
 - Click to open it

Warning: If you don't see the workspace, check that you're signed in with your university account, not a personal Microsoft account.
3. **Open the CloudRevenue Report**
 - Locate the report named **CloudRevenue - Sales Dashboard**
 - Click to open it in viewing mode
4. **Explore the Report Briefly**
 - Take 1-2 minutes to visually scan the report
 - Note what visuals are present (charts, tables, cards)
 - Don't analyze deeply yet - just get oriented

Part 2: Meet Copilot (5 minutes)

1. **Open the Copilot Pane**
 - Look for the **Copilot icon** in the top toolbar (sparkle/star icon)
 - Click it to open the Copilot side panel
 - The panel should appear on the right side of your screen

Note: The Copilot icon looks like a sparkle or star. If you don't see it, Copilot may not be enabled for your tenant — notify your instructor.

2. **Read the Welcome Message**
 - Copilot will introduce itself
 - Notice the suggested prompts it offers
3. **Try Your First Prompt**
 - Type: **Summarize this report**
 - Press Enter and wait for the response
 - Read the summary carefully

4. **Reflect** (CSAR: Reflect phase)
 - Did the summary capture the main points?
 - What did it miss?
 - Was anything surprising?
 - *What question would you ask next to clarify?*

Part 3: Practice the CSAR Loop (15 minutes)

Now you'll practice **dialog engineering** by applying the CSAR loop to business questions from the BRD.

CSAR in Action: Don't just fire off prompts. For each question set: 1. **Clarify:** What exactly do you need to know? 2. **Summarize:** Frame your question clearly for Copilot 3. **Act:** Ask and review the response 4. **Reflect:** Did it answer your need? What's missing?

Question Set 1: Revenue Performance **Clarify first:** What does the VP need to know about revenue? (Review the BRD)

Prompt 1: "What is the total revenue for 2024?"

Prompt 2: "Which product category generates the most revenue?"

Prompt 3: "Show me revenue by region"

Reflect: Did Copilot's answers match what you expected from the BRD requirements?

Your Notes: _____

Question Set 2: Trends and Changes

Prompt 4: "How has revenue changed over the last 6 months?"

Prompt 5: "Which quarter had the highest revenue growth?"

Prompt 6: "What's the trend for Enterprise customers?"

Your Notes: _____

Question Set 3: Product Analysis

Prompt 7: "Compare E3 and E5 license performance"

Prompt 8: "Which products have the highest average contract value?"

Prompt 9: "What products are growing fastest in APAC?"

Your Notes: _____

Question Set 4: Your Own Questions Based on the BRD, ask **3 additional questions** of your own:

Prompt 10: " _____ "

Prompt 11: " _____ "

Prompt 12: " _____ "

Part 4: Test Copilot’s Limits (5 minutes)

Try these prompts to explore where Copilot struggles:

"Predict next quarter's revenue"

"Why did revenue drop in March?"

"What should the VP do about churn?"

Warning: Copilot cannot predict the future or make business recommendations. It summarizes and analyzes existing data — it doesn’t replace executive judgment.

Reflection Questions (CSAR: Reflect): - Which questions did Copilot answer well? - Which questions did it struggle with? - What types of questions seem outside its capabilities? - *How does this inform your appropriate reliance calibration?*

Part 5: Practice Appropriate Reliance (5 minutes)

Now apply the **Appropriate Reliance Framework** from dialog engineering:

1. **Pick one insight from Copilot**
 - Example: “APAC region shows 15% growth”
2. **Ask the three questions:**

Question	Your Answer
Is the cost of an error high?	(If this number goes into a board presentation: YES)
Does Copilot have a track record here?	(First time asking: NO track record yet)
Can I verify quickly?	(Check the visual: YES, takes 30 seconds)

3. **Verify manually**
 - Find the visual that shows this data
 - Does the number match what Copilot said?
 - Is any context missing?
4. **Document the comparison**
 - Was Copilot accurate?
 - Did manual inspection reveal anything Copilot missed?
 - *Update your mental model: is Copilot reliable for this type of question?*

Building Trust Over Time: Each verification builds your calibration. After several labs, you’ll know which question types you can trust and which need checking. This is how dialog engineering develops from “verify everything” to “appropriate reliance.”

Lab Deliverable

Submit: A 1-page reflection document answering these questions:

Reflection Questions

1. **When did Copilot excel?**
 - List 2-3 types of questions where Copilot provided valuable, accurate answers
 - What made these questions effective?
2. **When did Copilot struggle?**
 - List 2-3 types of questions where Copilot’s answers were incomplete or unhelpful

- Why do you think it struggled?
3. **Verification Insights**
 - Describe your manual verification experience
 - Was Copilot’s answer accurate? What was missing?
 4. **Business Application**
 - As a business analyst, how would you use Copilot in your daily work?
 - What role should human judgment play?

Submission Format

- **File Name:** Lab1_Reflection_[YourName].docx or .pdf
- **Length:** 1 page (approximately 300-400 words)
- **Submit via:** Canvas → Assignments → Lab 1

Completion Checklist

Before submitting, confirm you have:

- Asked at least 10 natural language questions
- Tested at least 3 questions where Copilot struggled
- Manually verified at least 1 Copilot insight using appropriate reliance
- Completed all 4 reflection questions
- Submitted via Canvas

Key Takeaways

1. **Dialog Engineering > Prompt Engineering:** Structured conversation (CSAR) beats one-shot prompting
2. **The CSAR Loop:** Clarify → Summarize → Act → Reflect — use it for every AI interaction
3. **Copilot excels at:** Summarization, simple aggregations, trend descriptions, filtering data
4. **Copilot struggles with:** Predictions, causal explanations, strategic recommendations, data not in the model
5. **Appropriate Reliance:** Trust calibrated to track record — verify high-stakes outputs until confidence is earned
6. **Critical skill:** Knowing when to trust AI outputs vs. when to verify manually

Preparation for Next Lab

- Review your Prompting Journal entry
- Read: Microsoft Learn - Write Copilot Prompts for Power BI
- Think about: What report would you create for the VP of Cloud Sales?

Further Reading

For a deeper exploration of dialog engineering concepts introduced in this lab:

Correa, F., with Finch, A. (2026). *Dialog Engineering: The Cognitive Science of Human-AI Partnership*. Independently published. ISBN: 9798252564784.

“*Dialog engineering gives AI a partner — and that partner is you.*”