PREDICT·X

Guide to *Prompt Engineering* for T&E Management

Empowering T&E Professionals with Intelligent Prompting

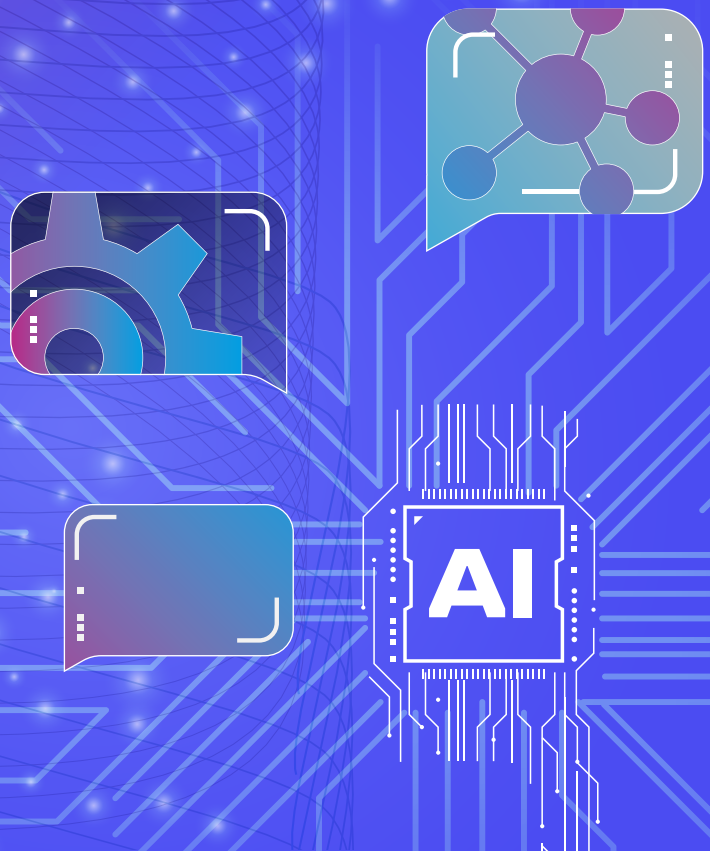

Contents

Introduction:

Navigating the AI Frontier in Travel & Expense Management

In today's rapidly evolving business landscape, **Artificial Intelligence (AI)** is no longer a futuristic concept but a practical tool transforming various industries, including Travel & Expense (T&E) management. **Large Language Models (LLMs)**, the advanced AI systems capable of understanding and generating human-like text, are at the forefront of this transformation.

Understanding how to effectively communicate with these intelligent systems is paramount for T&E professionals. This guide introduces the fundamental principles of **"Prompt Engineering"** – the art and science of crafting optimal inputs (prompts) to achieve desired outputs from LLMs. By mastering prompt engineering, you can unlock the full potential of AI tools, turning complex T&E data into actionable insights and streamlining your operations.

This e-book will equip you with the knowledge and techniques to interact with LLMs more effectively, ensuring you get precise, relevant, and valuable information every time. All examples provided are designed to be **100% educational** and demonstrate **practical applications within a T&E context using PredictX's Cogent AI**.

Chapter 1:

Understanding Large Language Models (LLMs) in T&E

What are LLMs?

Large Language Models (LLMs) are sophisticated AI programs trained on vast amounts of text data. This training allows them to understand, generate, summarize, translate, and even reason with human language. They learn patterns, grammar, facts, and context from the data, enabling them to respond to a wide range of prompts.

Why Prompt Engineering Matters for LLMs in T&E

For Generative AI: It's how you control the quality, tone, and format of the content you create.
For Agentic AI: It's the foundational skill. An AI Agent relies on a series of well-engineered internal prompts to create plans, execute tasks, and solve problems autonomously.

How LLMs Assist in T&E Management

LLMs can revolutionize T&E management by:

- **Automating Data Analysis:**
  Quickly sifting through vast datasets of travel bookings, expense reports, and supplier invoices.
- **Generating Reports:**
  Creating customized reports on spending patterns, compliance, and budget adherence.
- **Identifying Trends & Anomalies:**
  Spotting unusual spending, potential policy violations, or opportunities for cost savings.
- **Providing Insights:**
  Offering data-driven recommendations for supplier negotiations, policy adjustments, and travel optimization.
- **Answering Queries:**
  Responding to specific questions about T&E data in natural language.

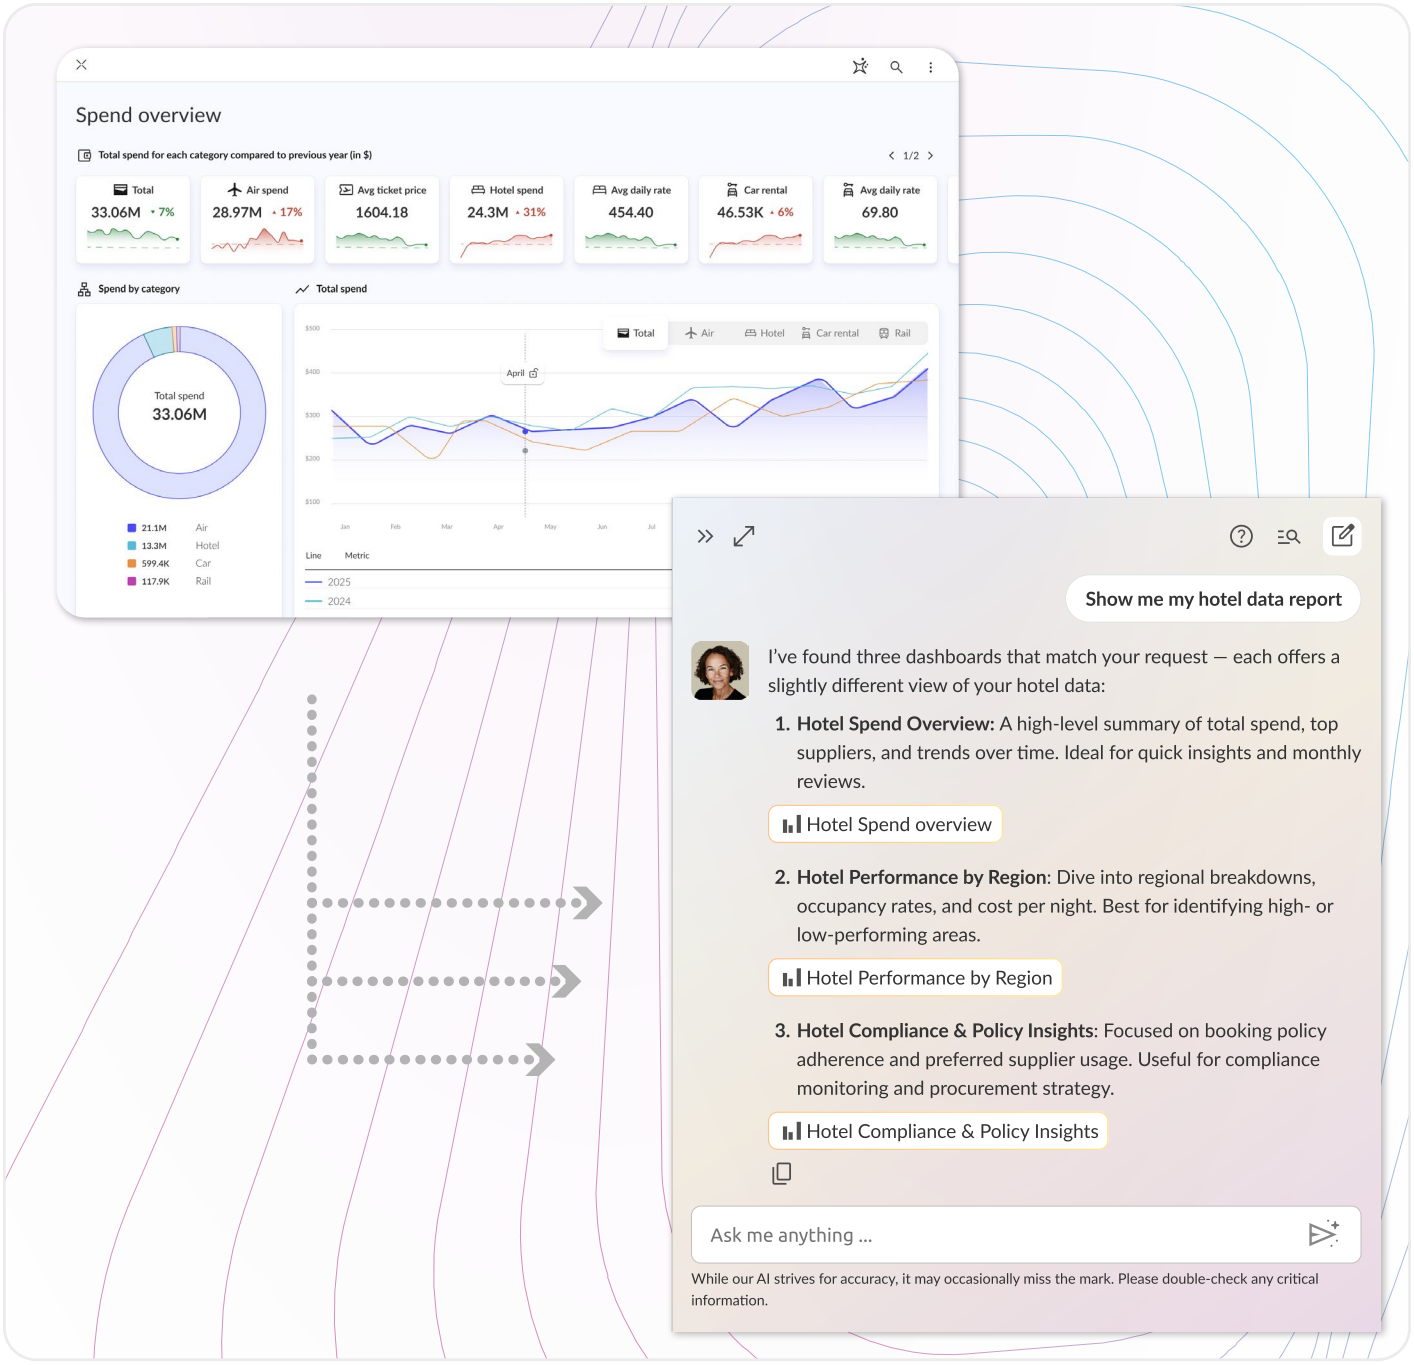

PredictX's **Cogent AI** is an example of an advanced AI solution leveraging these LLM capabilities to provide tailored insights and automation for T&E professionals. However, the quality of any LLM's output directly depends on the quality of your input. [This is where prompt engineering becomes crucial.]

Chapter 2:

Core Principles of Effective Prompting for LLMs in T&E

To get the best results from any LLM, your prompts need to be **clear**, **specific**, and **provide sufficient context**.

Principle 1: Clarity and Specificity

Vague prompts lead to vague responses. Be as precise as possible about what you want the LLM to do and what information it should use.

Ineffective Prompt: *"Show me travel costs."*

Why it's ineffective: An LLM doesn't know ***which*** travel costs (air, hotel, car?), ***when*** (last month, last year, specific quarter?), or ***how you want them presented*** (total, breakdown, by department?).

Effective Prompt: *"Provide a detailed breakdown of Q3 2025 air travel expenses for the EMEA region, grouped by airline, showing total spend and number of flights for each."*

Explanation: This prompt is highly specific:

What: "Detailed breakdown of air travel expenses."
When: "Q3 2025."
Where: "EMEA region."
How: "Grouped by airline, showing total spend and number of flights."

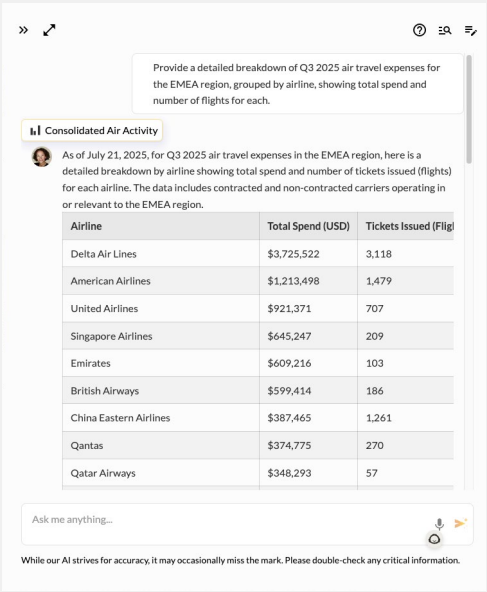

Principle 2: Providing Context

Context gives the LLM the necessary background information to understand your request fully and generate relevant responses. Think about what the LLM needs to know to perform the task accurately.

Ineffective Prompt: *"Analyze hotel spend for compliance."*

Why it's ineffective: An LLM doesn't know ***which*** compliance rules to apply, ***which period*** to analyze, or ***what specific aspects*** of hotel spend are relevant.

Effective Prompt: *"Considering our new travel policy effective July 1st, 2025, which caps hotel stays at $200/night for domestic travel and $300/night for international travel, analyze Q3 2025 hotel spend for compliance. Highlight any instances where the policy was exceeded and by what percentage."*

Explanation: This prompt provides crucial context:

Policy Details: Specifies the new policy, effective date, and specific caps for domestic and international travel.
Scope: Narrows the analysis to "Q3 2025 hotel spend."
Desired Output: Clearly states to "highlight any instances where the policy was exceeded and by what percentage."

Principle 3: Role-Playing and Persona

Instructing an LLM to adopt a specific role or persona can significantly influence the tone, style, and focus of its response, making it more tailored to your needs.

Ineffective Prompt: *"Give me ideas to save money on travel."*

Why it's ineffective: This is too generic. An LLM might provide general tips that aren't specific to your company's data or T&E context.

Effective Prompt: *"Act as a cost optimization consultant for a global enterprise. Based on our historical T&E data from the last fiscal year, identify the top three actionable areas where we can reduce overall travel and expense spend without significantly impacting employee productivity. For each area, provide a brief rationale and suggest a concrete strategy."*

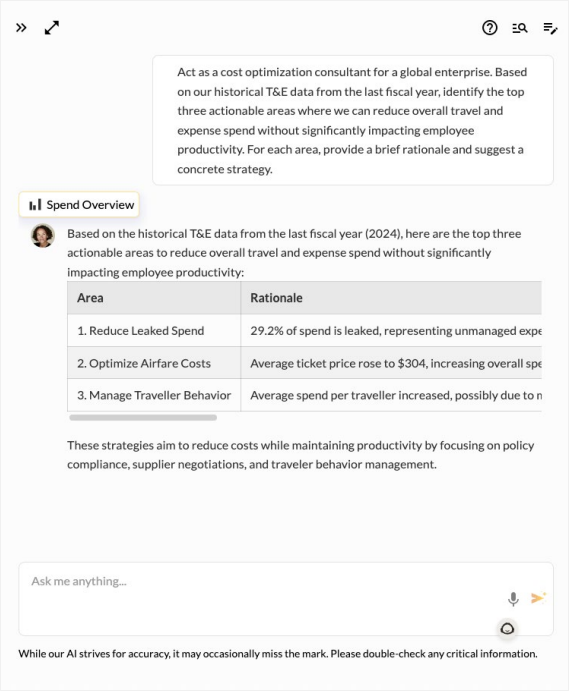

Explanation: This prompt assigns a persona:

Role: "Cost optimization consultant for a global enterprise."
Focus: Directs the LLM to analyze "historical T&E data" and provide "actionable areas" with "rationale" and "concrete strategies," aligning with a consultant's output.

Additional Example: *"As an air contracts manager, analyze the Star Alliance air contract, give me an overall summary using tables and give me savings opportunities."*

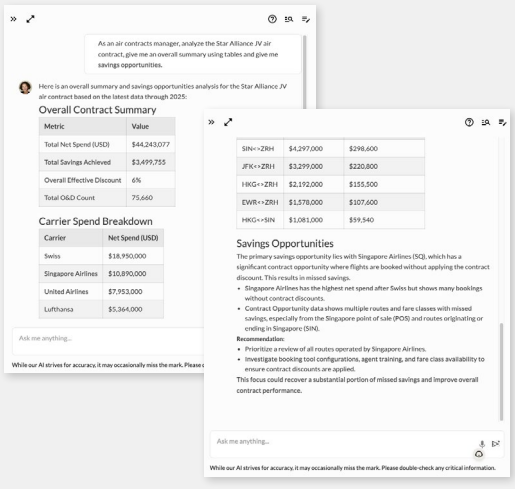

Explanation: This example specifically directs Cogent AI to take on the persona of an "air contracts manager" to provide a summary and identify savings opportunities related to a specific contract, using tables for clarity.

Chapter 3:

Essential Prompting Techniques for Maximizing LLM Potential in T&E

Beyond the core principles, specific techniques can further enhance your interaction with LLMs.

Technique 1: Zero-Shot Prompting

Zero-shot prompting is when you ask an AI to handle a task without giving it any prior examples, training, or detailed instructions. The AI relies on what it already knows to provide an answer immediately.

To craft an effective zero-shot prompt, consider these elements:

Subject (optional): State the main topic or focus of the task (e.g., "Subject: Research of public tech companies"). Use if the subject matter could be confusing or very specialized.

Action Verb: Use a clear command to direct the AI (e.g., "Analyze," "Summarize," "List").

Task Description: Specify what needs to be done with the subject (e.g., "Provide key financial metrics").

Context (Optional): Add details like audience, purpose, or scope (e.g., "For a business analyst preparing a report").

Tone/Format: Define the tone (e.g., professional) and output format (e.g., bullet points, table).

Constraints (Optional): Include limits like word count or specific requirements (e.g., "Focus on 2024 data, under 200 words").

Educational Example using Cogent AI:

"Provide a comprehensive overview of total travel spend, transaction volumes, and traveler counts over the past two years for a finance manager preparing a quarterly corporate travel budget review. In a professional tone, present data in a markdown table with clear year-by-year comparisons. Focus on data from 2024 to 2025, include percentage changes year-over-year, and highlight trends in spend and transactions."

Expected Cogent AI Response:

"Our total T&E spend for November 2024 was $X,XXX,XXX." (Assuming it has access to this data.)

Explanation:

This prompt combines all the elements mentioned above to guide Cogent AI to produce a specific, formatted report based on its existing knowledge.

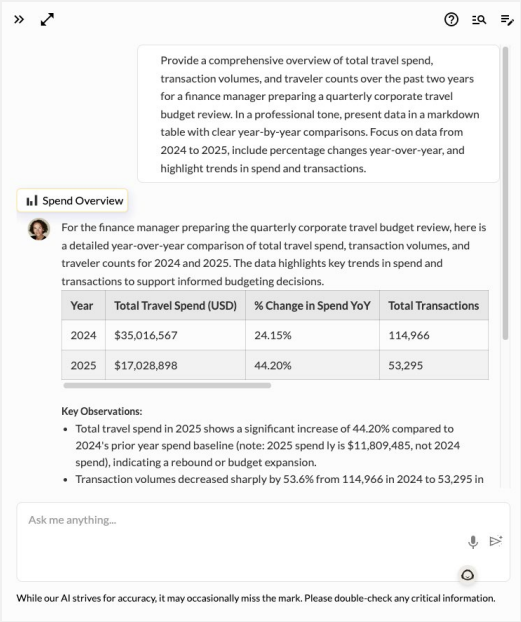

Technique 2: Few-Shot Prompting

When an LLM needs to understand a specific format, style, or interpretation, providing a few examples within your prompt can significantly improve the quality and consistency of its output.

Educational Example using Cogent AI:

"I need to categorize some miscellaneous expenses. Here are examples of how I want expense categories summarized:

- *`Meals & Entertainment` -> `F&B`*
- *`Airfare` -> `Flights`*
- *`Ground Transportation` -> `Local Travel`*

Now, apply this summarization format to the following list of expenses:

- *`Hotel Accommodation`*
- *`Conference Registration Fees`*
- *`Office Supplies`*
- *`Client Dinners`*
- *`Taxi Fare`"*

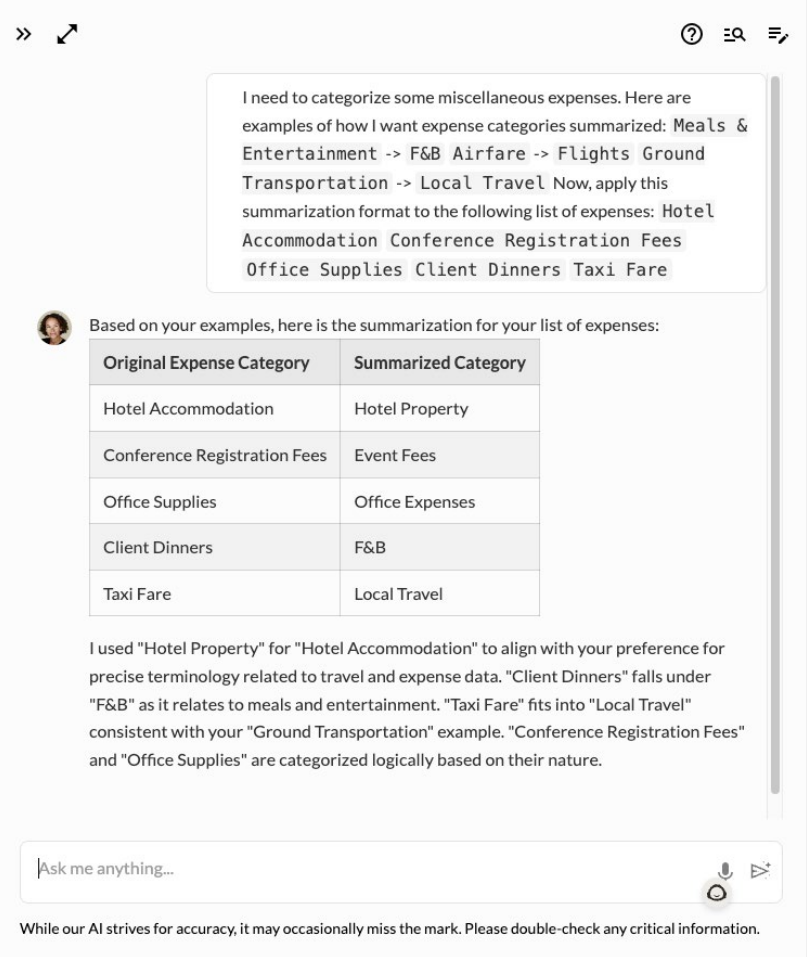

Expected Cogent AI Response:

- **`Hotel Accommodation` -> `Accommodation`**
- **`Conference Registration Fees` -> `Training/Events`**
- **`Office Supplies` -> `Office Ops`**
- **`Client Dinners` -> `F&B`**
- **`Taxi Fare` -> `Local Travel`**

Explanation:

By providing examples, you "teach" the LLM the desired mapping or summarization logic, even for categories it hasn't seen in the examples.

Technique 3: Chain-of-Thought Prompting

For complex tasks that require multiple steps of reasoning, guiding an LLM through a "chain of thought" can lead to more accurate and robust results. This involves breaking down the problem into smaller, sequential steps.

Educational Example using Cogent AI:

"I need to identify and summarize our top non-compliant travel expenses for Q1 2025. Please follow these steps:

1. *First, identify all employee travel expenses from Q1 2025 that exceeded policy limits by more than 15%.*
2. *Second, for each identified instance, list the employee's name, the expense amount, the specific policy violated, and any reason provided by the employee.*
3. *Third, calculate the total overspend across all identified non-compliant expenses.*
4. *Finally, summarize the top 3 categories (e.g., hotel, airfare, meals) where these non-compliant expenses occurred most frequently."*

Expected Cogent AI Response (structured):

"Step 1: Identified Non-Compliant Expenses (Q1 2025 >15% over policy)

- **[List of individual non-compliant expenses]**

Step 2: Details of Each Instance

- **Employee: Jane Doe, Amount: $750 (Hotel), Policy Violated: Hotel Cap ($500), Reason: Last-minute booking due to emergency.**
- **... (more instances)**

Step 3: Total Overspend

- **Total overspend for Q1 2025: $X,XXX.XX**

Step 4: Top 3 Non-Compliant Categories

1. **Hotel Accommodation (45% of overspend instances)**
2. **Business Class Airfare (30% of overspend instances)**
3. **Meals & Entertainment (20% of overspend instances)"**

Explanation:

This technique forces the LLM to process the request logically, reducing the chance of errors in complex analysis.

Technique 4: Using Delimiters

Delimiters are special characters (like triple backticks ```, quotes “”, or XML tags) used to clearly separate instructions from the input text that an LLM needs to process. This prevents misinterpretation.

Educational Example using Cogent AI:

“Summarize the following travel report, focusing on identifying any unusual spending patterns or significant deviations from budget. The report is enclosed within triple backticks:

```
[Start of Travel Report Text]
Travel Report: Q2 2025 Corporate Travel Summary
Department: Sales
Total Airfare: $150,000 (Budget: $120,000) - Note: 25% increase due to urgent client visits.
Total Hotel: $80,000 (Budget: $75,000)
Total Meals: $40,000 (Budget: $30,000) - Note: High spend on client entertainment in NYC.
... [More report details] ...
[End of Travel Report Text]
```

Expected Cogent AI Response:

“The Q2 2025 Sales Department travel report shows significant deviations from budget in airfare (25% over), and meals (33% over). The airfare increase is attributed to urgent client visits, and high meal spend to client entertainment in NYC. Hotel spend was slightly over budget. Further investigation into the specific instances of overspend for airfare and meals may be warranted.”

Explanation:

The backticks clearly tell the LLM which part of the prompt is the instruction and which is the data to be summarized.

Technique 5: Instruction Tuning / Output Constraints

You can refine an LLM's output by providing specific instructions on the desired format, length, or content.

Educational Example using Cogent AI:

"Generate a report on Q4 2024 car rental spend for the North America region. Ensure the report includes the following columns: `Vendor`, `Number of Rentals`, `Total Cost`, and `Average Daily Rate (ADR)`. Exclude any individual rentals under $50. Present the data in a markdown table format."

Expected Cogent AI Response (Markdown Table):

Vendor	Number of Rentals	Total Cost	Average Daily Rate (ADR)
Enterprise	120	$25,000	$70.00
Hertz	95	$20,000	$70.00
Avis	80	$18,000	$75.00
Budget	60	$12,000	$65.00
Total	**355**	**$75,000**	**$70.42**

Explanation:

This prompt specifies:

- **Content:** "Q4 2024 car rental spend, North America."
- **Columns:** Explicitly lists the required columns, including "Average Daily Rate (ADR)".
- **Filter:** "Exclude any individual rentals under $50."
- **Format:** "Present the data in a markdown table format."

Technique 6: Prompt Stacking

Prompt stacking is like giving an AI a list of instructions to follow one after another, like a recipe. You tell the AI exactly what to do, step by step, in a specific order, and it follows those steps to give you an answer. Each step builds on the one before it, like stacking blocks to build a tower.

Educational Example using Cogent AI:

"Using the hotel property spend data from 2022 to 2025, first provide a year-over-year summary of total spend and room nights. Then, analyze how the average negotiated rate and average daily rate have changed over these years. Finally, calculate the average spend per traveler for each year and identify which year had the highest efficiency in spend per traveler relative to room nights. Please present all findings in a clear table format and summarize key insights."

Expected Cogent AI Response (structured table and summary):

Year	Total Spend	Room Nights	Avg Negotiated Rate	Avg Daily Rate	Avg Spend per Traveler	Efficiency (Spend/Room Night)
2022	$X,XXX,XXX	YYY,YYY	$A.AA	$B.BB	$C.CC	D.DD
2023	$X,XXX,XXX	YYY,YYY	$A.AA	$B.BB	$C.CC	D.DD
2024	$X,XXX,XXX	YYY,YYY	$A.AA	$B.BB	$C.CC	D.DD
2025	$X,XXX,XXX	YYY,YYY	$A.AA	$B.BB	$C.CC	D.DD

Export to Sheets

Key Insights:

- **[Summary of year-over-year trends in total spend and room nights]**
- **[Analysis of changes in negotiated vs. daily rates]**
- **[Identification of the year with highest efficiency in spend per traveler relative to room nights, with explanation]**

Explanation:

This prompt clearly outlines a sequence of analytical steps for Cogent AI to follow, ensuring a comprehensive and structured output.

Chapter 4:

Advanced Prompting Strategies for LLMs in T&E

As you become more comfortable, you can employ advanced strategies to tackle more complex T&E challenges with LLMs.

Technique 1: Zero-Shot Prompting

This involves a conversational approach where you refine your prompts based on an LLM's previous responses. It's like having a dialogue to progressively narrow down or expand on an analysis. This is also known as Multi-turn Prompting.

Educational Scenario using Cogent AI:

User Prompt 1:
"Cogent, give me an overview of our company's top 5 travel spend categories for the last fiscal year."

Cogent AI Response 1:
"Our top 5 travel spend categories for the last fiscal year were:

1. *Airfare: $X,XXX,XXX*
2. *Hotel Accommodation: $Y,YYY,YYY*
3. *Meals & Entertainment: $Z,ZZZ,ZZZ*
4. *Car Rental: $A,AAA,AAA*
5. *Conference & Training Fees: $B,BBB,BBB"*

User Prompt 2 (Iterating on Response 1):
"Thanks, Cogent. Now, specifically for 'Airfare,' can you break down the spend by airline for the last fiscal year and highlight any airlines where our spend increased by more than 10% compared to the previous fiscal year?"

Cogent AI Response 2:
"Certainly. Here's the breakdown of airfare spend by airline for the last fiscal year, with year-over-year increases highlighted:

- *Airline A: $X1,XXX (15% increase from previous year)*
- *Airline B: $Y1,YYY (5% increase)*
- *Airline C: $Z1,ZZZ (20% increase from previous year)*
- *... (more airlines) ..."*

User Prompt 3 (Further Iteration):
"Cogent, for Airline A and Airline C, can you identify the top three routes that contributed most to the increased spend?"

Cogent AI Response 3:
"For Airline A, the top three routes contributing to increased spend were: [Route 1, Route 2, Route 3]. For Airline C, they were: [Route 4, Route 5, Route 6]."

Explanation: This iterative process allows you to conduct deep-dive analyses, starting broad and then progressively focusing on specific areas of interest.

Strategy 2: Tree of Thought Prompting

Tree of Thought prompting is like having a conversation with an AI where it thinks step-by-step to solve a problem, like you'd talk through ideas with a colleague. Instead of just giving one quick answer, the AI creates a "thought tree" by breaking the problem into smaller parts, like branches. It explores different ideas, checks which ones make sense, and builds on them in a logical, related conversation to find the best solution.

Educational Example using Cogent AI:

Initial Exploration Prompt:
"Can you provide an overview of the key air travel compliance KPIs for 2024, including compliance rates and estimated missed savings?"

This prompt initiates a broad summary to understand the main performance indicators and their financial impact.

Focused Deep Dive Prompt:
"Which air travel compliance KPI category in 2024 has the highest estimated missed savings, and what does this imply for potential cost savings?"

This prompt narrows the focus to identify the biggest opportunity area, encouraging reasoning about business implications.

Route-Level Analysis Prompt:
"For the KPI category with the highest missed savings, can you provide detailed route-level data including lost savings, compliance rates, and ticket counts for 2024?"

This prompt drills down to granular data, supporting detailed analysis and targeted action planning.

Explanation:

This strategy guides Cogent AI through a multi-stage reasoning process, allowing for more complex problem-solving and detailed analysis by breaking down the query into logical steps.

Strategy 3: Implementing Guardrails and Constraints

You can set explicit boundaries or rules for an LLM's responses to ensure they adhere to specific company policies, data privacy, or analytical frameworks.

Educational Example using Cogent AI:

Prompt:

"When analyzing supplier performance for our travel program, only use data from our list of preferred vendors. Do not suggest changes to our preferred vendor list unless explicitly requested and only based on objective performance metrics like on-time performance and cost-effectiveness. Furthermore, ensure any financial figures are presented in USD and rounded to two decimal places."

Explanation:

This prompt establishes clear guardrails:

- ***Data Source Constraint:*** *"only use data from our list of preferred vendors."*
- ***Action Constraint:*** *"Do not suggest changes to our preferred vendor list unless explicitly requested..."*
- ***Formatting Constraint:*** *"ensure any financial figures are presented in USD and rounded to two decimal places."*

Chapter 5:

Best Practices for Prompt Engineering in T&E

To consistently achieve optimal results from LLMs in a T&E context, integrate these best practices into your workflow.

1: Start Simple, Then Iterate:

Begin with a straightforward prompt. If the initial response isn't exactly what you need, refine your prompt by adding more detail, context, or constraints.

2: Be Clear and Concise:

Avoid jargon where plain language will suffice. Every word in your prompt should serve a purpose.

3: Provide Examples (Few-Shot):

When the output format or specific interpretation is critical, show the LLM what you expect with a few examples.

4: Define the Output Format:

If you need data in a specific structure (e.g., a table, a list, a summary paragraph), explicitly state it in your prompt.

5: Test and Refine:

Prompt engineering is an iterative process. Don't be afraid to experiment with different phrasings and techniques until you achieve the desired outcome.

6: Understand LLM Capabilities and Limitations:

While powerful, LLMs operate on the data they have access to and their training. Ensure your requests align with the type of T&E data and analytical capabilities available through the specific LLM you are using (like Cogent AI).

7: Use Delimiters for Clarity:

Especially when providing long inputs or multiple pieces of information, use delimiters to help the LLM distinguish instructions from content.

8: Leverage Persona/Role-Playing:

Guide the LLM to adopt a specific perspective (e.g., "Act as a financial auditor," "As a travel manager...") to get responses tailored to that viewpoint.

Conclusion:

Mastering Your T&E Data with Intelligent Prompting

Prompt engineering is a powerful skill that transforms how you interact with AI tools. By applying the principles and techniques outlined in this guide, T&E professionals can move beyond basic queries to conduct sophisticated analyses, generate insightful reports, and proactively manage travel and expense spend.

The ability to effectively communicate with AI is becoming as crucial as understanding financial spreadsheets or travel policies. Embrace prompt engineering, and unlock a new level of efficiency, insight, and control over your organization's T&E management. PredictX's Cogent AI is your intelligent assistant; prompt engineering is the language that empowers it to serve you best.

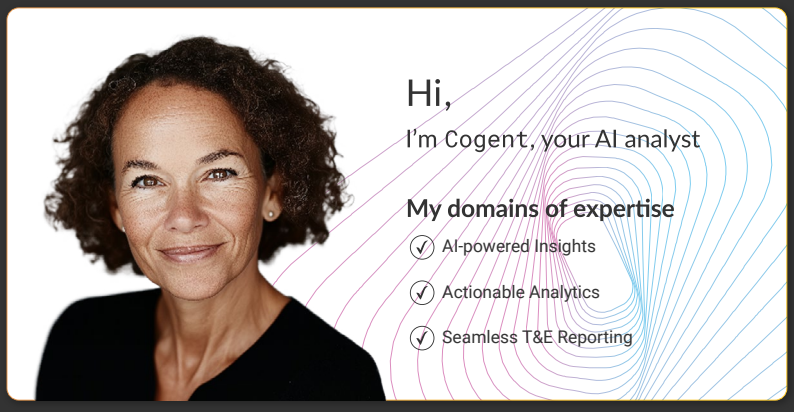

The future of travel and expense is about to change,

PredictX is here to lead the way.

Request a trial today: predictx.com/solutions/cogent
Questions? Contact us at: sales@predictx.com