POWERMIND
BY QUERIXA

PowerMind Sample Report: Practical Application Guide

To demonstrate the practical application of the PowerMind visual, we have designed a sample report utilizing a straightforward dataset. This report intentionally includes complex data relationships to highlight how PowerMind can navigate—and overcome—standard Power BI design limitations.

Visual Placement & Layout Ideas

The sample report features three distinct implementations of the PowerMind visual to inspire your dashboard designs:

- **Embedded:** Integrated alongside other charts for contextual insights.
- **Full Page:** Taking up the entire canvas for a dedicated, immersive chat experience.
- **Pop-up Style:** Acting as an overlay or modal for on-demand assistance without crowding the dashboard.

Handling Data: Raw Data vs. Measurements

Instead of feeding pre-calculated DAX measures into PowerMind, this sample report demonstrates the power of utilizing raw data.

- We add raw data directly to the visual's field parameters.
- We then use the **Conversation Guidance > Assistant** Instructions to explain, in plain language, how the agent should calculate KPIs (e.g., revenues, transactions, profit).
- Once the data fields are added, the agent builds a "context table." We highly recommend toggling to the default table view and downloading this data to inspect exactly what context the AI is analyzing.

Managing Filter Interactions

Like any standard Power BI visual, PowerMind interacts with page filters. However, to provide macro-level insights and broad comparisons, the AI sometimes requires an unfiltered view of the dataset.

- **Best Practice:** In this sample, we modified the visual interactions to disable cross-filtering from certain slicers. This ensures the AI retains the "whole picture" of the data when users ask comparative questions across different categories.

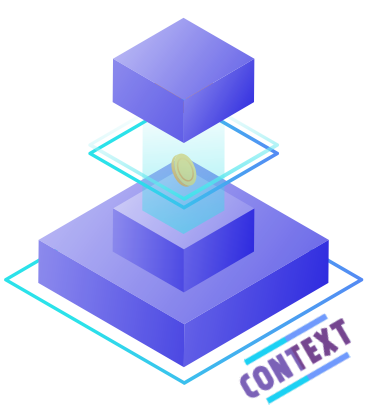

Overcoming Row Context Limitations

Power BI's underlying architecture relies on row context, which dictates how the context table is aggregated before the AI processes it. This can lead to calculation discrepancies if not handled correctly:

- **The Average Dilemma:** If you add an average DAX measure to the visual, the context table displays the average per row. If a user asks the AI for the "global average," the AI might average those row-level averages, leading to an incorrect mathematical result.
  - *Solution:* Provide the "count of rows" alongside the average, and instruct the agent to recalculate the true average using a weighted calculation.
- **The Count Dilemma:** In our sample, "transactions" represents the count of rows in the fact table. However, adding fields to the visual creates an aggregated context table, meaning it will have fewer rows than the raw fact table. If the AI simply counts the rows of the context table, the transaction count will be wrong.
  - *Solution:* We added a "Distinct Count of Sales ID" to the visual and explicitly instructed the agent to sum this column to find the true number of transactions.

By understanding the row context principle, you can easily guide the PowerMind agent to provide accurate, highly advanced insights every time.