The AI Strategy - Growth Compass

The AI Strategy: One Model or Multiple?

We need Multiple Models working in a pipeline. A single model like Gemini/GPT/mistral cannot know our specific sales trends without being fed our data via specialized models.

A. Custom ML Models

We must train these ourselves because they are specific to our business geography and product lifecycle.

- Demand Forecasting Model (Prophet or XGBoost): Predicts SKU-level demand by region to reduce stockouts.
- Churn/Risk Model (Random Forest): Identifies retailers whose ordering frequency is dropping.
- Productivity Scoring (Clustering): Groups sales reps based on Call-to-Order efficiency rather than just total volume.

B. Generative AI (The Reasoning Agent)

We do not need to train Gemini or GPT. Instead, use them via API with an Agentic approach.

- Why skip RAG for numbers? RAG is for text. For our sales targets, we need Text-to-SQL.
- How it works: We give the LLM our Table Schema (list of columns). When we ask "Why is Brand X falling behind?", the LLM writes a SQL query to get the data, analyzes the result, and writes the answer.

How to Train the Custom Models

For the 15-20% or more revenue impact, our Demand Forecasting model is key. Here is the training flow:

1. Data Preparation: Create a 3-year or more historical window of our DMS sales.
2. Feature Engineering: Add external shifters like holidays, paydays, and weather.
3. Training: Use an 80/20 split. Train on 80% of the past data; test if the model can predict the other 20% accurately.
4. Deployment: Run the model every weekend (Friday) night so sales reps have fresh Daily Targets every Saturday morning.

Key Components for Test & Learn Goal

To allow to run experiments, implement A/B Testing Analytics in the dashboard:

- The Sandbox: Assign Region A as the test group and Region B as the control.
- Statistical Significance: Use a simple ML script to calculate if the growth in Region A was real or just noise.

Summary Checklist: What We Need

- Data Warehouse: BigQuery or Snowflake or any database to hold the huge data.
- ML Pipeline: Python (Scikit-learn/XGBoost) for Forecasting and Churn.
- LLM API: Gemini/GPT (for natural language insights and SQL generation).
- BI Tool: User interface for the visual dashboard.

The Growth Compass Architecture

Link : [Growth Compass Architecture]