Triple Whale

The Ultimate LLM Prompting Guide for Ecommerce Brands

Introduction

Just by clicking this link, you gave us some data. Every website, every application, and any interaction with an ecommerce brand becomes a data point that a brand can use — if it knows how.

But the amount of data collected doesn't always correlate with value, and many brands are turning to generative AI and LLMs (large language models) — like ChatGPT and Claude — to turn that data into actionable intelligence. LLMs can help analyze data quickly and adapt to unique business requests, but there are specific guidelines an ecommerce brand should follow in order to obtain the best results possible with the power of generative AI.

In this article, we'll explore:

- What LLMs can do for your ecommerce business
- How to prepare data from your ecommerce store for analysis
- Guidelines for crafting effective prompts that generate actionable insights
- How to solve common ecommerce challenges with LLMs

Finally, we'll introduce you to the next generation of ecommerce intelligence with agentic AI, and the Moby Agents that are trained on ecommerce data and ready to uncover powerful insights more quickly, with no data collection, preparation, or prompting required.

A comprehensive breakdown of what agentic AI is and how AI agents can help you grow and scale your ecommerce business is here: The Ultimate AI Agents Playbook for Ecommerce Brands.

Table of Contents

What LLMs can do for your ecommerce business

Remember how we said most brands are just drowning in data?
We consider LLMs to be a bit of a life raft. An LLM can take the raw data you plug into it and uncover some insights you may not have seen initially (or ever) because you weren't sure what you were looking at.

There are several applications for LLMs to optimize an ecommerce brand, but to name a few:

- **Product search optimization.** Conventional search is based on keyword matching, but LLM-powered search can interpret natural language to deliver results that are based on context and intent.
- **Conversational chatbots.** Rule-based chatbots of the past followed rigid conversation paths (highly annoying), but LLM-powered chatbots understand context, remember the conversation history, and can deliver a truly personalized experience.
- **Personalized product recommendations.** LLM-powered recommendations understand the relationship between products and a nuanced understanding of a customer's preference across multiple touchpoints, to deliver suggestions that make sense for that individual.
- **Fraud detection and prevention.** From payment fraud to refund scams, you're one wrong click away from getting defrauded. LLMs have much improved pattern recognition abilities, and are capable of uncovering fraud attempts more efficiently than before.
- **Customer feedback analysis.** Modern LLMs can interpret customer emotions expressed across product reviews, social media comments, support tickets, and chat interactions altogether.
- **Ecommerce data analysis.** LLMs can take huge amounts of data from ecommerce platforms, ad platforms, inventory management platforms and more to distill into actionable insights.

What LLMs like ChatGPT and Claude cannot do

While they're powerful, LLMs can't access your data directly. You'll need to export data from various platforms, prepare it in a way the LLM can understand, and craft a prompt that will provide the response you need. Because these tools don't have access or real-time connections to your Shopify store, advertising accounts, Google Analytics, and more, they are limited in the insights they can provide.

The quality of the insights depends on the quality and completeness of the data you provide to the LLM.

Here are the dos and don'ts of LLM prompting:

DO

- ✓ **Be specific about data format**
  Mention if it's a CSV, PDF, etc. and describe the key fields, like SKU, revenue, ROAS
- ✓ **Provide clear business objectives**
  Eg. "Find products with declining sales" or "Identify pricing anomalies above $X"
- ✓ **Include relevant timeframes**
  Specify date ranges, seasonal periods, or comparison windows
- ✓ **Define metrics clearly**
  Explain how you calculate conversion rate, AOV, or profit margins
- ✓ **Specify the output format**
  Request tables, charts, summary reports, or actionable recommendations
- ✓ **Provide business context**
  Share industry, target market, or competitive landscape when relevant
- ✓ **Ask for specific analyses**
  Customer segmentation, inventory optimization, etc.

DONT

- × **Assume ChatGPT/Claude knows your data structure**
  Don't just say "analyze this data" without any context
- × **Ask vague questions**
  "What do you think about this data?" "Make it better."
- × **Mix different time periods**
  Without noting when data is from different quarters or years, you will run into problems
- × **Use internal jargon**
  Don't use company-specific terms or KPIs that the LLM doesn't understand
- × **Leave the output format open-ended**
  - LLMs like clear instructions
  - Open-ended output requests lead to generic analyses that aren't useful
- × **Assume universal ecommerce rules**
  B2B vs. B2C are different, and should be analyzed in distinct ways.
- × **Request everything at once**
  Break complex analyses into focused questions

DO	DONT
✓ **Include data quality notes** Mention known issues, missing fields, or data collection changes	✕ **Hide data limitations** This can lead to incorrect conclusions
✓ **Request actionable insights** “What should we do next?” “Which products need attention?”	✕ **Start at observations** Raw findings without business implications aren’t useful
✓ **Use sample data first** Test with a subset before analyzing large datasets	✕ **Upload massive files immediately** Without understanding processing limitations

LLMs also can’t make any business decisions for you, nor can they guarantee the accuracy of the predictions they make. They work best as partners to help you understand your data and to generate hypotheses, but the strategic decisions will remain in your hands.

Unlike ChatGPT and Claude, Triple Whale’s Moby Chat Need to know something about your data? Just ask. Try Moby Chat to get instant business-specific insights from your always-on smart assistant. Moby has direct access to your data, and is basically an ecommerce-trained generative AI tool that already understands the nuances of ecommerce.

The foundation is the **Triple Whale Data Platform,** which combines all of the ecommerce data you have available to pull all of your metrics into a single source of truth. It is also layered with sophisticated attribution models, forecasting abilities, and the business intelligence tools to analyze your data from every angle.

Moby is the only AI that understands ecommerce, and is able to revolutionize the way you access and analyze your data.

> **Need to know something about your data?**
> Just ask. Try Moby Chat to get instant business-specific insights from your always-on smart assistant.

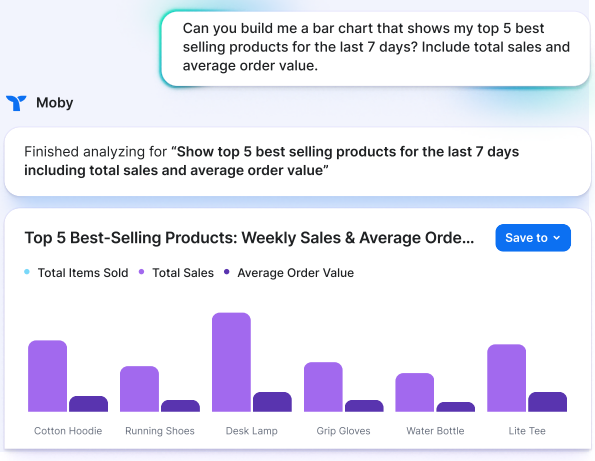

LLM prompting fundamentals

To be successful with LLMs, you need to understand how they process information and structure your requests in a way that most effectively produces your desired end result.

To provide valuable insights, the LLM needs extremely clear directions.

Here are some general guidelines:

1. **Give the LLM context.** LLMs perform best when they understand not only what you're asking, but why you're asking it. Let the LLM know the context of your request by outlining your business model, some current challenges, goals you have for the analyses, and any constraints that might affect the recommendations.
2. **Specificity is a necessity.** If you provide a vague request, you'll get a vague answer. Being more specific with what you need will generate a more valuable output.
3. **Structure the request.** A well-structured prompt produces better results.
4. **Iterate and refine.** A broad first analysis can help identify patterns, and follow-up questions will allow you to drill down into a specific area of interest. Ask for clarification when recommendations aren't clear or when initial findings raise new questions.

When in doubt about how to formulate a prompt, remember that it should be measurable, obtainable, bounded, and yielding. We like to use the acronym **MOBY:**

M

Measurable (include identifiable metrics)
Instead of: *"Tell me about sales trends."*
→ **Try:** "Analyze our Q4 2024 online sales data to identify trends in customer behavior across different regions, using order revenue and net profit."

O

Obtainable (within the AI's capabilities)
Instead of: *"Predict our total revenue for the next decade."*
→ **Try:** "Using historical monthly revenue data, forecast the next three months' revenue trends."

B

Bounded (Time-bound with clear scope)
Instead of: *"Give me an overview of our performance."*
→ **Try:** "Summarize our performance metrics for the past six months, focusing on customer retention and engagement trends."

Y

Yielding (Relevant and specific goals)
Instead of: "What's going on with our website performance?"
→ **Try:** "Evaluate our website performance by comparing conversion rates, bounce rates, and average session duration between the first and last month of Q4 2024."

How to prepare data for LLM analysis

The quality of the analysis an LLM delivers is directly connected to the quality of the data it is provided with.

The data preparation phase is crucial, and it will represent the difference between receiving generic observations and actionable intelligence to move your business forward. Below are some guidelines for collecting, organizing, and getting data ready for LLM analysis.

Data collection & access

1. Start by identifying all key data sources relevant to your business

- Ecommerce platforms (Shopify, WooCommerce, BigCommerce)
- Analytics (GA4)
- Advertising platforms (Meta, Google Ads, TikTok, etc.)
- Email marketing tools (Klaviyo and similar)
- Customer service platforms
- Inventory management systems
- Payment processors

2. Set up proper data export capabilities to ensure you can access data in usable formats

- CSV exports from each platform
- API access where available
- Regular automated exports (weekly/monthly)
- Historical data going back at least 12 months

3. Before diving into analysis, confirm that your tracking infrastructure is correctly set up

- UTM parameters on all marketing campaigns
- Conversion tracking pixels are properly installed
- Customer lifetime value calculations
- Product performance metrics
- Attribution data across touchpoints

Data organization & preparation

1. Standardize your data format by creating consistent naming conventions

- Product categories and SKUs
- Campaign names across platforms
- Customer segments
- Date formats
- Currency and metric units

2. Group your data into logical, analysis-ready structures

- Customer data: demographics, purchase history, LTV
- Product data: sales, inventory, margins, categories
- Marketing data: campaign performance, attribution, ROI
- Financial data: revenue, costs, profitability
- Operational data: fulfillment, returns, customer service

3. Clean and validate data to ensure data quality

- Remove duplicates
- Handle missing values
- Verify data accuracy
- Normalize formats across sources

Get data ready for analysis

1. Build reusable templates for common analysis needs

- Monthly performance dashboards
- Product performance reports
- Customer segmentation analysis
- Marketing ROI assessments
- Inventory optimization reports

2. Include contextual information for AI analyses

When using tools like ChatGPT or Claude, always provide relevant context to improve analysis quality:

- Business goals and KPIs
- Seasonal patterns
- Market conditions
- Recent changes or campaigns
- Industry benchmarks (if known)

Effective prompting for ChatGPT/Claude

As evidenced by many prompts gone wrong or halfway to what you expected, crafting an effective prompt is both an art and a science. The best prompts combine clear business objectives with structured data requests and specific output requirements. If a prompt follows this approach, it is more likely to deliver insights that are actionable immediately, rather than generic observations that will require additional interpretation.

Since ecommerce analysis typically involves multiple variables and competing priorities, the prompting strategy should reflect the complexity of decision-making involved. Simple question-and-answer scenarios aren't often useful, and effective prompts will provide the generative AI model with enough structure to deliver valuable insights.

Here are some guidelines for how to prompt LLMs:

1. Craft specific analysis requests

Instead of: "analyze my data"

→ Use:

```
Analyze my Q4 sales data to identify which product categories had the strongest
growth, what customer segments drove the most revenue, and where I should focus
inventory investment for Q1.
```

2. Ask the right questions

The questions you ask an LLM make all the difference. For example:

- "Which customer segments have the highest LTV and how can I acquire more customers like them?"
- "What product combinations are most frequently bought together and how can I optimize my bundling strategy?"
- "Which marketing channels are most cost-effective for acquiring customers with above-average order values?"

3. Request actionable outputs

Always ask for:

- Specific recommendations with reasoning
- Prioritized action items
- Expected impact of changes
- Metrics to monitor success
- Implementation timelines

Best Practices for ChatGPT/Claude Prompting and Analyses

Start with clear objectives

Define what business decision you're trying to make

- Specify the timeframe for analysis
- Identify the key stakeholders who will use the insights

Provide complete context

- Include business model details
- Explain any unusual events during the data period
- Share your hypotheses or concerns
- Mention any constraints or limitations

Use iterative analysis

- Start with high-level analysis
- Drill down into specific area of interest
- Ask follow-up questions based on initial findings
- Request additional analysis as patterns emerge

Focus on actionability

- Always ask: "What should I do with this information?"
- Request prioritized recommendations
- Ask for implementation guidance
- Specify success metrics and monitoring plans

Examples for ChatGPT/Claude prompts for ecommerce analysis

The following examples demonstrate how to apply LLM analysis to a few common ecommerce challenges:

- Evaluating ad performance across platforms
- Inventory management and forecasting
- Customer insights
- Ad creative analysis
- How performance is pacing

Each example will include the business context, data preparation requirements, and some detailed prompts you can adapt for your specific needs. These are templates and not scripts to be followed exactly. As each business has unique characteristics, these should be reflected in how you structure your analysis requests and interpret the results. Adapt, iterate, and refine the prompts to get the exact output you're hoping to achieve.

1. Evaluate ad performance across platforms

A modern ecommerce brand advertises on multiple platforms with budgets spread across Facebook, Google Ads, TikTok, and more – which can make it difficult to decipher how effective any one ad (or any one platform) is for the brand.

A brand needs to know:

Are my ads being run efficiently?

It's not just about cost-per-click or number of impressions, but instead the brand needs to determine if the ad spend it generating meaningful business outcomes at a sustainable cost. Beyond individual campaign performance, it's important to understand how different platforms complement each other in the overall marketing funnel.

Are some platforms better than others for acquiring new customers?

While one platform may drive high awareness, another might be better for converting customers. Understanding how the overall funnel works will allow a brand to budget funds for platforms based on specific business needs.

Should we change how we budget ad spend based on performance?

A dynamic budget allocation can respond to real performance data rather than gut feelings or outdated assumptions. But, in order to accomplish that, a brand must understand both their current performance and historical trends in order to make decisions about where to decrease or maintain spending.

How to collect and organize ad performance data for LLM analysis

To get a comprehensive view of advertising performance across platforms, a brand must build a system that compiles all of the essential metrics from each platform in the same place:

- **Metrics to pull from each platform.** Ad spend, impressions, reach, clicks, CTR, conversion rates, conversion value, ROAS, CPC, CPM, and CPA
- **Pull metrics from across specific timeframes.** For example, if you are interested in comparing recent trends with historical long-term performance, you can pull both 30-day and 365-day periods.
- **Data organization.** Create a master spreadsheet with consistent naming conventions across all platforms, including campaign names, audience labels, and metric definitions.
- **Create comparison views.** Build side-by-side platform comparisons that show performance metrics in the same format, and include percent changes over time to quickly identify the top and
  bottom performers.

Once all of the data is organized, the LLM will be able to transform the raw data into insights that might take hours to identify manually.

Ad performance LLM prompt

Here is a suggested LLM prompt for evaluating ad performance across data platform:

```
# Ad Performance Data Analysis Request

## Context I have comprehensive advertising performance data across multiple

platforms for both 30-day and 365-day periods. I need actionable insights to optimize ad
spend allocation and improve overall performance.

## Your Role
Act as a senior digital marketing analyst with expertise in performance marketing,
budget optimization, and cross-platform advertising strategy.

## Analysis Framework Please conduct a systematic analysis following this structure:

### Performance Efficiency Assessment
- Calculate and rank platforms by ROAS (highest to lowest)
- Identify cost per acquisition leaders and laggards
- Evaluate cost efficiency for impressions and clicks
- Flag platforms with spending inefficiencies

### Customer Acquisition Analysis
- Segment new vs. repeat customer acquisition by platform
- Calculate true cost per new customer by platform
- Assess customer lifetime value by acquisition source

### Budget Optimization Strategy
- Recommend budget reallocation based on performance data
- Identify platforms for increased/decreased investment
- Suggest optimal budget percentage distribution
- Identify spending caps due to diminishing returns

### Trend & Pattern Recognition
- Compare 365-day vs. 30-day performance shifts
- Identify platforms with positive/negative momentum
- Note seasonal patterns affecting performance
- Highlight concerning metric declines

## Required Outputs

### Format Requirements
1. Executive Summary (3-4 key findings + primary budget recommendation)
2. Platform Scorecard (Letter grades A-F with 2-3 bullet justifications each)
3. Budget Allocation Table (Specific percentages/amounts with current vs. recommended)
4. Action Priority Matrix (Top 5 actions ranked by impact/urgency)

### Analysis Standards
- Support all recommendations with specific data points
- Include confidence levels for major recommendations (High/Medium/Low)
- Provide reasoning for each letter grade assignment

- Quantify potential impact where possible (e.g., "could improve ROAS by X%")

## Data Interpretation Guidelines
- Prioritize statistical significance over absolute numbers for smaller platforms
- Consider platform maturity when evaluating performance
- Factor in attribution windows and measurement differences
- Account for seasonal business patterns in trend analysis  Please begin your analysis
once I provide the data.
```

As you can see, this is a very detailed and organized prompt that requires accurately compiled data to produce useful results.

Here's what the output of that exact prompt would look like in Moby Chat, which automatically accesses your brand's data, no uploading required.

With Triple Whale, all of the data is already compiled accurately for this analysis, and there is even an agent available to automatically run analyses to evaluate ad performance across platforms, and it’s called the Channel Budget Allocation Agent.

The Channel Budget Allocation Agent provides real-time, optimal budget distribution across marketing channels. It is able to analyze performance data and spending efficiency over the past 30 and 365 days.

Focusing on metrics like spend, ROAS, and conversion values, the Agent can help a brand maximize returns and adjust recommendations for budget allocation based on current performance as well as historical trends.

How the **Channel Budget Allocation Agent** works:

- **Performance analysis.** Continuously monitors campaign performance metrics and spending efficiency across all channels at once.
- **Real-time optimization.** Provides dynamic budget allocation recommendations based on current performance as well as historical trends.
- **Goal alignment.** Adjusts recommendations based on your specific business objectives and revenue targets.
- **Spend monitoring.** Tracks budget utilization and alerts you to opportunities for reallocation or expansion.
- **Impact forecasting.** Projects expected results from recommended budget changes to inform decision-making.

→ Want to skip all of the data collection, organization, and prompt crafting required to do this analysis in an LLM yourself?

Give the **Channel Budget Allocation Agent** a try today.

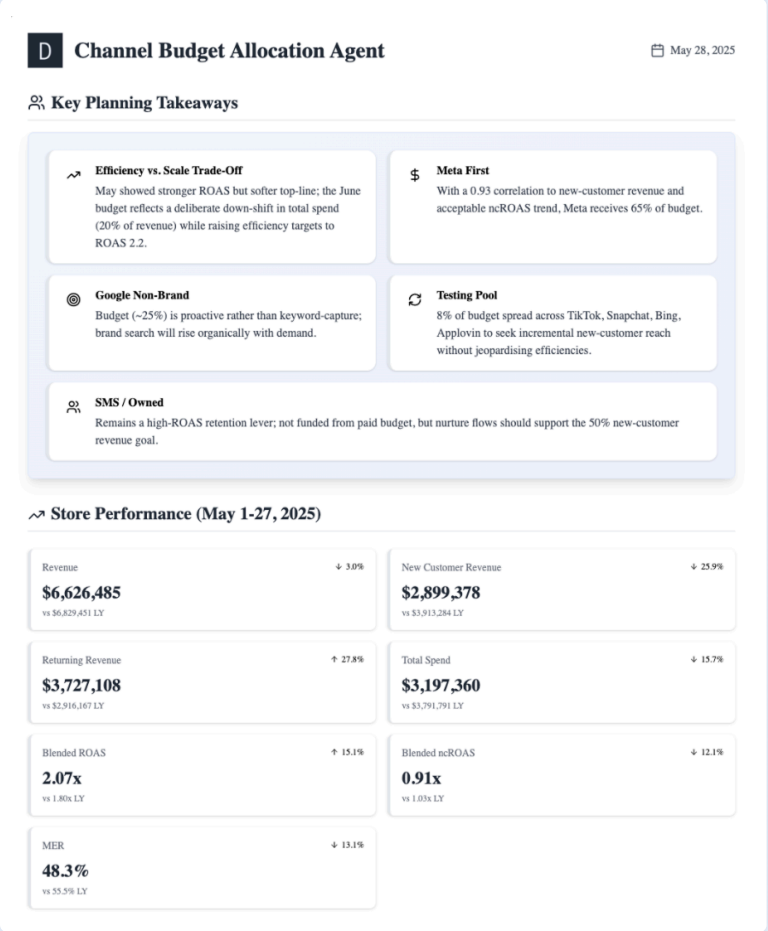

2. Monitor and optimize inventory trends

Proper inventory management is essential for an ecommerce brand, because an overstock of products ties up cash flow and hikes up storage costs, and too little inventory means missed sales and frustrated customers. With thousands of SKUs across multiple categories, each product could even have different demand patterns, seasonality, and profit margins.

A brand needs to know:

Which products are performing best, and which are underperforming?
Understanding top sellers versus slow movers should also consider profit margins, inventory turnover, and storage costs. A high-revenue product that sits in a warehouse for months might be less valuable than a lower-revenue item that turns over weekly.

Are there any unusual patterns in sales and inventory data?
Sudden sales spikes could indicate viral social mentions, stockouts from competitors, or the early stages of a trending product category. Being able to identify these patterns early allows you to capitalize on opportunities or prepare for demand shifts ahead of time.

Can we predict future sales trends confidently?
Accurate forecasting is a cornerstone of efficient inventory management. This means understanding not just the historical patterns but how external factors like seasonality, marketing campaigns, and market trends influence demand for each product.

How to collect and organize inventory data for LLM analysis:
To effectively prepare your inventory data for LLM-powered analysis, start by collecting and organizing the following key categories:

Inventory Data Collection & Organization Checklist for LLM Analysis

Product Identifiers and Classifications

Capture detailed product-level data for granular performance tracking:

- ☐ SKU or Product ID
- ☐ Product name
- ☐ Product variants (track individually if performance differs)
- ☐ Product attributes (e.g., price, seasonality, target demographic)

Current Inventory Metrics

Track real-time stock and operational data:

- ☐ Stock levels across all warehouses and locations
- ☐ Reorder points and minimum stock thresholds per SKU
- ☐ Cost of Goods Sold (COGS)
- ☐ Lead times for restocking
- ☐ Per-unit storage costs

Sales Performance Data

Monitor performance trends over time:

- ☐ Daily, weekly, and monthly sales quantities and revenue
- ☐ Gross profit margins
- ☐ Discount frequency
- ☐ Return rates
- ☐ Customer Acquisition Cost (CAC)
- ☐ Lifetime Value (LTV) by product category

Time-Based and Operational Metrics

Evaluate efficiency and inventory flow:

- ☐ Days of supply remaining at current sales pace
- ☐ Inventory turnover rates (compared to industry benchmarks)
- ☐ Average days to sell from stocking to purchase
- ☐ Stockout frequency and duration per SKU

Seasonal and External Factors

Incorporate broader context into analysis:

- ☐ Impact of holidays and promotional periods on sales
- ☐ Historical demand trends
- ☐ Seasonal fluctuations

Organizing your data for analysis

Once you've collected the necessary data, the next step is to structure it in a way that enables efficient, reliable analysis — especially when working with LLMs:

1. Start by identifying all key data sources relevant to your business

Gather and merge key inventory-related documents:

- Full product-level sales reports
- Current inventory status reports
- Cost analysis reports

2. Create comprehensive margin calculations

Build detailed profitability models by including all cost components in your spreadsheet:

- COGS
- Shipping to warehouse
- Storage fees
- Payment processing and platform fees

3. Structure data for AI analysis

Format your dataset in a machine-readable way:

- Use a CSV file where each row represents a unique product or variant
- Include separate columns for all relevant metrics
- Ensure data points align with consistent time intervals

4. Ensure historical depth

Provide at least 12 months of data to identify trends, account for seasonality, and enable more accurate forecasting.

Inventory management LLM prompt

Once you've successfully collected and organized the data, an LLM can provide a sophisticated analysis that would require extensive manual work to achieve otherwise.

Here is a suggested LLM prompt for effectively monitoring and optimizing inventory levels:

```
# Ecommerce Inventory & Sales Analysis

I've uploaded our inventory and sales data. Please analyze and provide:

## 1. Performance Analysis
- Identify top and bottom 20% of products by:
- Profitability
- Turnover rate
- Total contribution
- Flag high-revenue/low-profit products (and vice versa)

## 2. Demand Forecasting
- Predict next month's sales by product category
- Identify stockout risks for next 30-60 days

## 3. Inventory Optimization
- Recommend optimal stock levels for top 50 SKUs
- Consider lead times and service level targets
- Identify overstocked items for liquidation

## 4. Seasonal Insights
- Analyze seasonal patterns
- Recommend inventory build-up timing for peak periods
- Identify shifting seasonal trends

## 5. Action Items (Next 30 Days)
Prioritize by financial impact with confidence levels:
- Products to reorder immediately
- Products to reduce orders for
- Pricing/promotional strategies to improve turnover
```

Similar to the ad performance analysis, the inventory analysis can also be accomplished automatically with one of Triple Whale's Agents: the **Product & Inventory Pulse Agent.**

The **Product & Inventory Pulse Agent** provides up-to-the-minute data on inventory levels and product performance by synthesizing data from various sources to provide key insights about stock levels, product performance, and demand forecasting. It can quickly identify unusual patterns that could impact your bottom line, and provides tailored suggestions to optimize stock management and improve sales.

How the **Product & Inventory Pulse Agent** works:

- **Data integration.** The agent consolidates data from various platforms, including sales, inventory, and market trends for a comprehensive view.
- **Performance monitoring.** Continuously tracks product performance metrics, identifying top sellers and underperformers.
- **Trend analysis.** Utilizes advanced algorithms to detect trends and patterns in sales and inventory data.
- **Anomaly detection.** Flags unusual inventory movements or sales spikes, allowing for timely intervention.
- **Insight delivery.** Provides daily insights and recommendations directly to your dashboard to ensure you never miss critical information.

→ Skip all of the inventory data collection and organization you'd need to complete yourself and instead click 'run' on the **Product & Inventory Pulse Agent!**

D **Product & Inventory Pulse**

TL;DR

Critical inventory issues detected with 10 products at risk of stockout and 8 showing negative inventory counts. Top performers include The Perfect Neckline Bra with strong revenue but concerning sell-through rates. Activewear shows exceptional profitability (55% contribution margin) while several products have fulfillment costs exceeding 50% of revenue, severely impacting margins.

Inventory Health Overview

⚠ **At-Risk Products**

Product	Inventory	USPW	Weeks on Hand
Comfort Shaping Bra (Black/4X)	-2	20.8	-0.1
Adjustable Comfort Bra (Blue/XL)	-7	40.4	-0.2
Racerback Sports Bra (Pink/L)	1	70.7	0.0
Smoothing Brief (White/M)	8	140.6	0.1
No Strings Attached Bra (Sand/M)	31	50.9	0.6

% Sell-Through Analysis

Product	7-Day	30-Day	60-Day
Shaping Lift Cami (White/L)	250.0%	102.6%	101.5%
Neckline Pullover Bralette (Black/L)	100.0%	100.0%	100.0%
Perfect Neckline Bra (Scarlet/L)	100.0%	100.0%	100.0%
Smoothing Brief (White/3X)	31.0%	55.4%	73.9%

Top Revenue Generators - Inventory Position

3. Understand and utilize customer value data to boost LTV

Customer lifetime value (LTV) is essential for businesses to stay profitable, yet many businesses don't keep their finger on that metric as much as they might AOV or conversion rate. LTV is arguably one of the most important metrics to understand what is driving the most long-term value for customers, and can help brands optimize everything from product development to how they allocate their ad spend.

A brand needs to know:

What drives customers to purchase certain products, and which products correlate with high LTV?
Some products serve as a gateway purchase that leads to repeat business, while others might just be a one-and-done.

Which characteristics define their highest-value customers?
Beyond demographics, brands need to understand behavioral patterns, acquisition channels, engagement preferences, and timing of purchases. High-LTV customers often share specific traits that can guide your targeting and retention strategies.

Where do their most valuable customers come from and how do they behave differently?
Different acquisition channels often produce customers with vastly different lifetime values. Understanding the differences can allow brands to allocate marketing budget toward the channels that tend to generate more sustainable growth.

How to collect and organize customer data for LLM analysis:
To ensure your customer data is optimized for AI-driven insights, use the following checklist to gather and structure all essential information categories.

Customer Data Collection Checklist for LLM Analysis

Ecommerce Platform Data

Export comprehensive customer purchase history and behavior:

- ☐ Frequency, recency, and monetary value (RFM) for each customer
- ☐ Average Order Value (AOV)
- ☐ Total lifetime spend
- ☐ First and most recent purchase dates
- ☐ Detailed category preferences
- ☐ Cart abandonment rates
- ☐ Browsing behavior
- ☐ Device usage
- ☐ Coupon utilization

Advertising and Acquisition Data

Unify acquisition data across channels with standardized tracking:

- ☐ UTM parameters and conversion API data
- ☐ Campaign associated with each customer acquisition
- ☐ First-click attribution data
- ☐ Customer Acquisition Cost (CAC) by channel
- ☐ Total number of touchpoints before purchase
- ☐ Messaging types that resonate with high-LTV customers

Email & SMS Marketing Engagement

Evaluate engagement effectiveness by customer segment:

- ☐ Open rates and click rates
- ☐ Post-email or post-text purchase behavior
- ☐ Segmentation by customer value tiers
- ☐ Correlation between engagement and customer lifetime value

Customer Feedback & Surveys

Capture sentiment and intent through structured feedback:

- ☐ Post-purchase surveys
- ☐ Net Promoter Score (NPS) surveys
- ☐ On-site feedback tools
- ☐ Key purchase motivations
- ☐ Product-level feedback
- ☐ Data on brand perception

Website Behavior Analytics

Use behavioral data to deepen customer journey insights:

- ☐ Pages visited before purchase
- ☐ Time spent on product pages
- ☐ Content engagement metrics
- ☐ Site search queries
- ☐ Behavioral heatmaps (e.g., Hotjar)
- ☐ Comparison between high-value and average customer behavior

Organizing your data for analysis

1. Create a unified customer database

Structure data as a CSV file with customer email or order ID as the unique identifier. Each row should represent one customer with columns for all collected metrics. This unified view enables AI to identify patterns across different data sources.

2. Establish customer value segments

Define meaningful customer tiers based on spending and engagement:

- Flag the top 10% of customers by lifetime value
- Create additional segments: top 25%, middle 50%, bottom 25%
- Include RFM scores to capture recency, frequency, and monetary value dimensions

3. Map the acquisition journey

Add key fields to reflect the acquisition path:

- First-touch and last-touch channel
- Total number of touchpoints
- Time elapsed from first interaction to purchase

4. Document behavioral patterns

Include behavioral columns to track recurring engagement traits:

- Purchase frequency trends
- Product or category preferences
- Responsiveness to promotions
- Interaction with specific content formats

Customer value data LLM prompt

Once you've successfully gathered and organized the data, it's time to prompt the LLM. Here's an example of a prompt to investigate customer lifetime value:

```
I've uploaded our customer database with purchase history, acquisition data, and
engagement metrics. Please analyze:

## 1. High-Value Customer Profile
- Identify characteristics of top 10% customers by LTV
- Compare their acquisition channels, behaviors, and preferences vs. average customers

## 2. LTV Prediction
- Predict which new customers will become high-value based on first 30-90 days
- Identify strongest early indicators of long-term value

## 3. Channel Performance
- Rank acquisition channels by average customer LTV
- Calculate LTV-to-CAC ratios by channel
- Compare customer behaviors by acquisition source

## 4. Retention Strategy
- Identify at-risk customers
- Recommend retention tactics by customer segment
- Highlight highest-success intervention points

## 5. Growth Recommendations Prioritize by revenue impact with confidence levels:
- Strategies to increase overall LTV
- Product bundling opportunities
- Communication optimization
- Acquisition channel improvements Segment all recommendations by implementation
difficulty.
```

There is, of course, a Triple Whale Agent that can do this a lot quicker: the **Customer Lifetime Value Analysis Agent.** This agent automatically synthesizes customer data into clear insights, helping you identify high-value segments, optimize engagement strategies, and drive retention growth. This agent can inform marketing efforts by understanding what resonates with your audience, and specifically which types of campaigns lead to improved engagement.

How the **Customer Lifetime Value Analysis Agent** works:

- **Data integration.** The agent consolidates customer data from various sources, providing a holistic view of customer interactions.
- **Behavior analysis.** Advanced algorithms analyze customer behavior patterns, identifying trends and anomalies.
- **Segmentation.** Customers are segmented based on value, preferences, and engagement levels, allowing for targeted marketing strategies.
- **Insight generation.** The agent generates clear, actionable insights and recommendations tailored to your business objectives.
- **Performance tracking.** Monitor the effectiveness of your strategies in real-time, adjusting as necessary to maximize impact.

→ Get details on your most valuable customers quickly and more easily with the **Customer Lifetime Value Analysis Agent.**

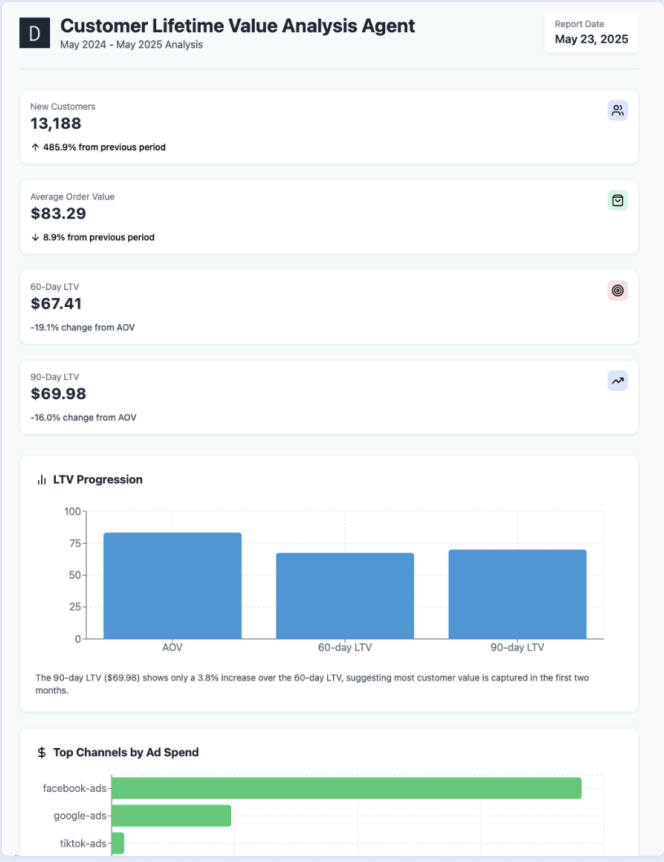

4. Analyze ad creative performance

Determining why certain ads succeed and why other ads fail is a task that would take several manual hours of analysis, which is what makes it a great application for LLM analysis. Deconstructing the visual, textual, and structural elements of a brand's creative can help them understand which components are resonating with their target audience to drive conversions.

The value for this type of analysis is that it combines quantitative performance data from the ad or ecommerce platforms with a qualitative creative assessment. A brand is always able to see which ads have the highest conversion rates, but LLMs can help them understand why certain ads succeed, and how to replicate that success systematically.

A brand needs to know:

Which ads are actually converting customers into buyers?
It's not enough to know which ads get clicks or even which generate the most impressions. Brands need to understand which specific creatives are driving actual conversions, and which are acquiring new customers versus just re-engaging existing ones.

What specific creative elements (images, copy, format) contribute to success?
High-performing ads often share common visual themes, messaging approaches, or structural elements that aren't immediately obvious. Understanding these patterns allows a brand to systematically replicate success rather than rely on guesswork for future creative development.

How can underperforming ads be improved or replaced with concepts from successful ads?
Every underperforming ad represents wasted budget and missed opportunities. Brands need strategies to either optimize existing weak performers or replace them with concepts based on proven successful elements from their best ads.

How to collect and organize ad creative data for LLM analysis:
To prepare your ad creative data for AI-powered analysis, use the checklist below to ensure you've captured both performance metrics and creative context across platforms.

Ad Creative Data Collection Checklist for LLM Analysis

Ad Performance Data

Export core data from ad platforms:

- ☐ Campaign structure details
  (campaign names, objectives, targeting parameters)
- ☐ Key performance metrics
  (impressions, CTR, CPC, conversion rates, ROAS)
- ☐ Indicators of new vs. returning customers
- ☐ Cost per acquisition and per-conversion values

Creative Elements

Collect the visual and textual components of your ads:

- ☐ Screenshots of all ad variations
  (mobile and desktop)
- ☐ Thumbnails and brief descriptions for video ads
- ☐ Complete ad copy elements
  (headlines, primary text, CTAs)
- ☐ Visual themes, color schemes, and overall design language

Exporting and Structuring Performance Data

- ☐ Export 30–90 days of campaign data from ad platforms in CSV format
- ☐ Include fields such as campaign names, ad names, targeting information, and performance metrics
- ☐ Prioritize recent data, while ensuring enough volume for pattern detection

Documenting Creative Assets

- ☐ Capture screenshots for each creative variant
- ☐ Record all text components and calls-to-action
- ☐ Annotate visual themes, color schemes, and design strategy
- ☐ For video ads, document key frames and summarize the content

Combining for LLM-Ready Analysis

- ☐ Create unique identifiers to link performance data with corresponding creative assets
- ☐ Add contextual information, including targeting criteria and ad placement
- ☐ Document any promotional offers, discounts, or campaign-specific variables that could influence performance

Creative ad performance LLM prompt

After all of the data is collected and organized, present it to the LLM with a comprehensive prompt as provided here:

```
# Ad Creative Performance Analysis
I've uploaded 60 days of performance data and creative assets for my advertising
campaigns. Please analyze to optimize creative strategy:

## Analysis Goals
- Identify high-converting creative elements
- Find patterns in top-performing ads to replicate
- Pinpoint weaknesses in underperforming ads
- Develop actionable creative recommendations

## Visual Analysis
- What visual themes, colors, and design elements drive performance?
- Which product photography styles convert best?
- How do visuals differ between top and bottom performers?
- What visual mistakes should I avoid?

## Copy Analysis - Which headline structures and messaging work best?
- What calls-to-action generate highest response rates?
- How does copy length affect performance by ad format?
- What tone and value propositions resonate most

 ## Format & Targeting Insights
- Which ad formats perform best by objective?
- How does performance vary by placement?
- Do creative preferences differ across audiences?
- What works best for new customer acquisition vs. retargeting?

## Deliverables
1. **Executive Summary** - Key findings and immediate actions
2. **Performance Comparison** - Side-by-side analysis of top vs. bottom performers
3. **Best Practices Guide** - Creative guidelines based on my data
4. **Improvement Recommendations** - Specific fixes for underperforming ads
5. **Creative Framework** - Template for evaluating future concepts

Include specific examples from my ads with reasoning for each recommendation.
```

The **Meta Creative Strategy Agent** can synthesize data to deliver comprehensive reports that highlight key performance metrics, identify trends, and provide actionable recommendations. It is able to transform raw data into clear, actionable reports to guide your advertising strategy, to save hours on manual reporting.

- **Data-driven insights.** Transform raw data into clear, actionable reports that guide your advertising strategy.
- **Efficient boost.** Save time on manual reporting and focus on creative strategy and execution.
- **Performance optimization.** Identify high-performing ads and eliminate underperformers to maximize ROI.

→ Identify and replicate the strategy of your highest-performing ads with the **Meta Creative Strategy Agent.**

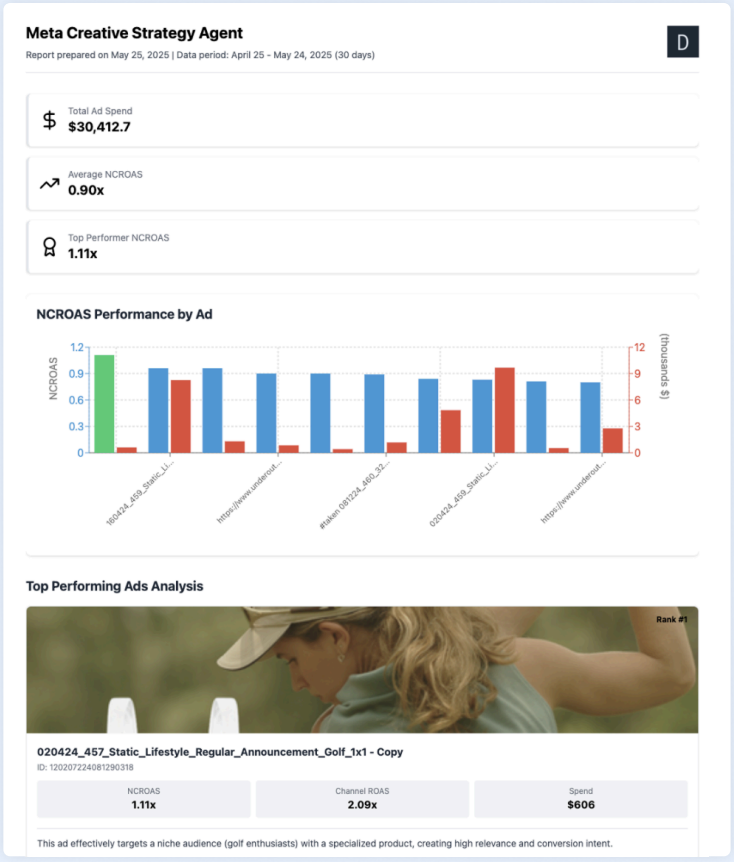

Bonus: Prompting Moby AI to analyze video ads

To get a detailed analysis of video ad performance, you can upload performance metrics downloaded from the ad platform that includes the following data:

- Average dropoff rate
- Average video views (total video views divided by number of ads)
- Average video duration
- Total ad spend on videos

When prompting **Moby AI**, be sure to:

1. **Be specific.** The structure of the report as well as the analysis requested should be clearly defined in the prompt, as well as the metrics used within it.
2. **Make the request obtainable.** Ensure you are asking for an analysi
3. s that is within the bounds of data that you've provided.
4. **Use Markdown in your "report" prompts.** Markdown is a language to structure and format text for use with AI models because it enhances their understanding. This will provide Moby with clearly differentiated sections of the report as well as specific requests.

Here is an example of a prompt formatted using Markdown that requests a report on video ad performance:

```
#Report Title: Best Performing Video Ad Analysis

##Sub-title: Use this analysis to make more winning video ads based on current best
performers

##Tertiary-title: Best performing video ads are determined by videos with the lowest
drop-off rate in the first 25 seconds.
-You should also provide the data ranges included in this report

#Your Directions You've been given data and analysis amongst my top video ads.

#Video Performance Metrics Start by providing me with an overview of my video ad
performance with the following aggregate metrics

-Average dropoff rate (first 25 seconds)
-Average video views (total video views divided by the number of ads)
-Average video duration
-Total spend on all video ads
```

Using the video ads with a "dropoff_0s_25" **designation of "low" or "average"**, compile the "video description", "hook description" and "first 25 second description" and summarize the common themes amongst the video ads with a low or average dropoff designation. You should use the video description and the ad copy for this analysis. Make it distinct in your report whether you're describing actions taking place in the video or ad copy.

Provide me with an analysis that describes the common themes from the video description that lead to it being a great video ad. You should bullet list out common themes found in the following parts of the video and any commonalities in the style of the video, actors being used, messaging, psychological principles, voice over, music, etc.

-full video description,
-the "hook" (first three second description),
-the "first 25 seconds" description

#Output Requirements
- Build the following sections using the data and insights captured in the previous step

1. Video Performance Metrics
2. Common Success Themes
3. Ads with Low or Average Dropoff Rates

-Display each of the video ads provided in a two-column card view where the **video url is displayed as a square image and is playable upon clicking**. The cards should be displayed side by side, next to eachother as two column and and multiple row matrix with two ads per row. The ad name and ad id should be displayed above the image with the rest of the metrics and descriptions below the video image in bullet format

4. Creative Brief
5. Ads with High Drop Off Rates -Display each of the video ads provided in a two-column card view where the **video url is displayed as a square image and is playable upon clicking**. The cards should be displayed side by side, next to each other as two column and and multiple row matrix with two ads per row. The ad name and ad id should be displayed above the image with the rest of the metrics and descriptions below the video image in bullet format

Remember that after the qualitative debrief, I want my video image cards displayed as a column x row format (2x) where I have two video image cards per row.

A prompt like this using **Moby AI** would provide a report that looks like this, and provides clear performance metrics for the video ads, common success themes, and playable files with the video content for easy viewing.

→ **Chat with Moby** to get a holistic and actionable look at your video ad performance.

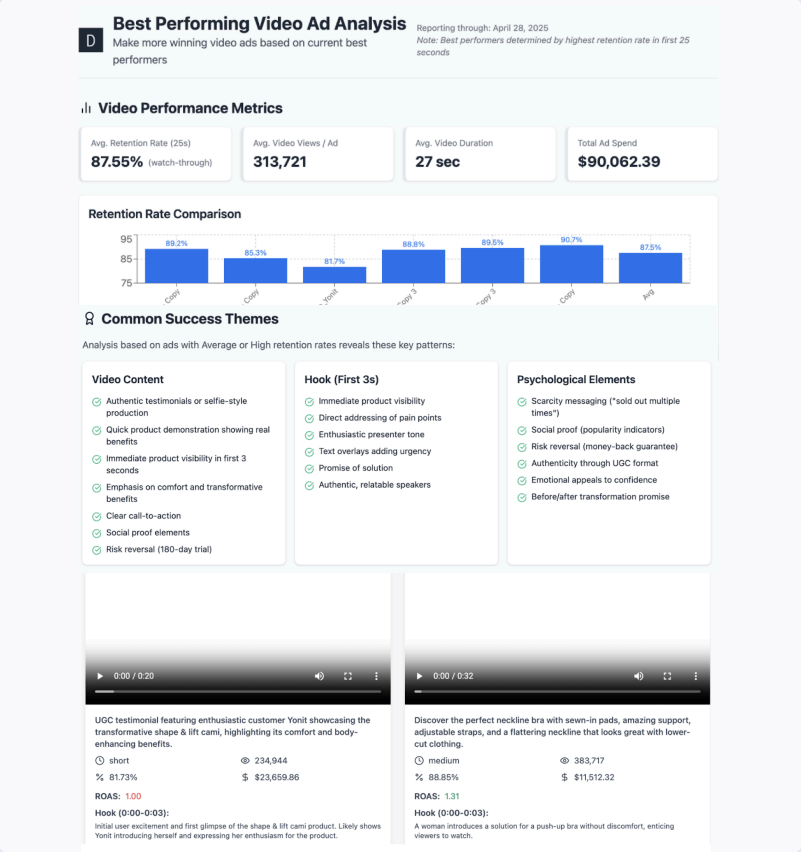

5. Ensure your performance is pacing in line with goals

It's easy to fall off course when you're operating a fast-paced ecommerce business.

Traditionally, brands would attempt to course correct halfway through the month, when the analyses are finally complete. By the time a brand realizes it's behind pace to meet month-end goals, it's too late to adjust advertising spend, inventory decisions, or promotional strategies. There's a critical need for real-time goal tracking, so you can make informed decisions more quickly.

To build a pacing analysis from the ground up, you need three core data components:

1. Current performance metrics
2. Historical comparison data
3. Clearly defined targets

Each component needs to be comprehensive and consistently formatted to enable meaningful analysis.

How to collect performance data for pacing analysis with an LLM:
To run a reliable pacing analysis with an LLM, you first need to collect a complete snapshot of your current performance, historical context, and clearly defined growth targets — all structured in a way that supports accurate comparisons and forecasting.
Essential data to collect:

1. Current performance metrics (month-to-date):

- **Revenue data.** Total MTD, by category, by channel, AOV, units sold
- **Customer data.** New customers acquired, returning purchases, CAC
- **Marketing performance.** Ad spend by channel, ROAS, traffic sources
- **Operational metrics.** Inventory levels, fulfillment times, return rates

2. Historical comparison data:

- Data from same month last year (complete month for context)
- Same days of month last year (true month-to-date comparison)
- Prior complete months for seasonal pattern analysis

3. Clearly defined growth targets:

- Overall revenue growth percentage
- Category-specific growth targets
- New customer acquisition goals
- Profit margin objectives
- Channel-specific performance expectations

Organizing your data for analysis

Once you've gathered the right performance, customer, and marketing data, the next step is organizing it in a way that enables reliable, fast-paced analysis. Here's how to structure your datasets and reporting environment so you can get the most out of LLM-powered insights.

Here’s how to structure your data for accurate pacing:

Organize Your Ecommerce Platform Data

Pull key performance metrics from platforms like Shopify, WooCommerce, or BigCommerce:

- Generate month-to-date (MTD) and year-over-year (YoY) comparison reports
- Export reports as CSV files
- Focus on:
  - Orders and revenue (both gross and net)
  - Average Order Value (AOV)
  - Product performance
  - Customer segmentation

Compile Marketing Performance Data

Aggregate MTD marketing metrics across your entire media mix:

- Export data from ad platforms (Meta, Google, TikTok, etc.)
- Include:
  - Spend
  - Impressions
  - Clicks
  - Conversions
  - Attributed revenue
- Don't forget to download:
  - Email marketing performance
  - Organic social metrics
  - Affiliate or partnership channel results

Align Your Analytics & Attribution Data

Use Google Analytics (or similar tools) to create custom reports that complement your media data:

- Compare current performance with the same period last year
- Focus on:
  - Traffic sources
  - Conversion rates
  - Revenue attribution
- Export this data in a format that can be easily matched with your advertising performance

Build An Analysis Workbook

Consolidate all of your structured data into a single workbook to enable faster, repeatable analysis. Suggested tabs include:

- **Executive Summary. Key metrics:** Last Year MTD, Current MTD, YoY Change, Monthly Target, Pacing to Target.
- **Daily Performance.** Daily breakdowns and trends across core KPIs.
- **Category Performance.** Product category-specific pacing and growth analysis.
- **Channel Performance.** Side-by-side comparisons of all marketing channels.
- **Targets & Forecasting.** Defined business goals and projections based on current pacing.

Goal-pacing LLM prompt

With all of this performance data and goals organized, you can request an analysis from the LLM using a prompt like this:

```
# Month-to-Date Performance Analysis

I've uploaded MTD performance data, YoY comparisons, and growth targets. Please analyze
pacing and provide recommendations:

## Performance Assessment
- Are we on track for monthly and annual revenue targets?
- Which metrics show strongest/weakest performance vs. targets?
- What percentage of monthly targets achieved vs. time elapsed?

## Channel & Category Breakdown
- Which marketing channels exceed/underperform expectations?
- Are product categories significantly ahead or behind target?
- How is customer acquisition pacing vs. annual goals?
- Which channels show best/worst ROI currently?

## Forecasting & Trends
- Based on daily trends, forecast end-of-month performance
- Identify concerning downward trends needing attention
- Highlight metrics with positive momentum to capitalize on
- Compare seasonal patterns to last year

## Strategic Recommendations
- What immediate actions improve underperforming areas?
- Where should we increase/decrease ad spend for monthly targets?
- Which promotional or inventory decisions could close gaps?
- How should we reallocate resources between channels?

## Deliverables
1. **Executive Dashboard** - Green/Yellow/Red pacing status
2. **End-of-Month Forecast** - With confidence intervals
3. **Priority Action Plan** - Ranked by impact potential
4. **Budget Reallocation** - Channel recommendations with spend suggestions
5. **Daily Monitoring Checklist** - For remainder of month
6. **Risk Assessment** - Biggest threats to targets

Include specific numbers, percentages, and reasoning for each recommendation.
```

Of course there's an agent that can do all of this automatically for you, and it's called the **Performance Pacing Agent.** It can quickly analyze current month-to-date performance against year-over-year growth targets and provide a comprehensive, data-driven report that empowers you to make informed decisions and optimize strategies for success.

How the **Performance Pacing Agent** works:

- **Data collection.** The agent automatically gathers month-to-date performance metrics and year-over-year growth targets for analysis.
- **Progress tracking.** Visual progress bars display your current performance against set targets, making it easy to assess your standing.
- **Pacing analysis.** The agent evaluates your current pace and forecasts whether you will meet your monthly targets by month-end.
- **Strategic recommendations.** Receive tailored insights and recommendations based on your performance data to optimize your strategies.
- **Visual forecasting.** Generate detailed visualizations that illustrate forecasted performance and target intersections for key metrics.

→ Want to make performance pacing a whole lot easier?
Try the **Performance Pacing Agent** today.

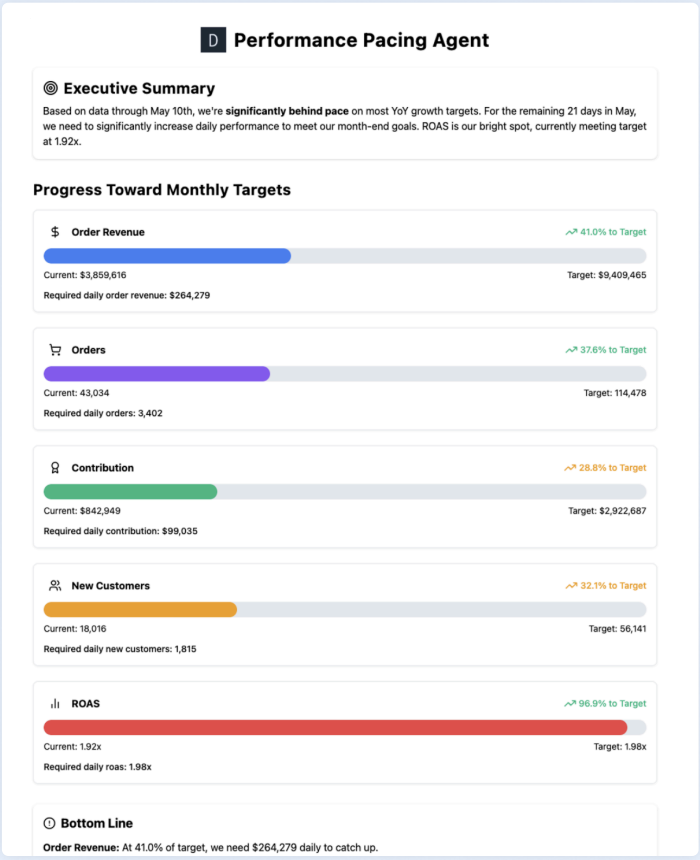

Moby Agents: the first ecommerce intelligence agents built for your entire brand

The above examples are only five of more than 100 dedicated agents that have already been developed by the Triple Whale team to specifically tackle the most common ecommerce data analysis requests. And we're developing new agents every day!

Moby Agents are ready to help you maximize performance at every stage of your business:

- **Media Buyer Collection.** Agents that focus on acquiring customers and optimizing ad strategies.
- **CRO Specialist Collection.** Agents designed to monitor every aspect of the business to drive conversions.
- **Operations Leader Collection.** These agents are focused on boosting profitability and scaling your brand.
- **Retention Marketer Collection.** Agents that optimize engagement strategies to keep your customers coming back.

Agentic AI is the latest innovation in artificial intelligence.

The Ultimate AI Agents Playbook for Ecommerce Brands will make sure you're up-to-speed with how AI agents can accelerate your growth!

Why Moby Agents are better than doing it yourself

In short, it saves you a ton of time and energy so you can focus on the more important things–like taking action on all of those intelligent insights. But, in more detail:

1. Data is correct, current, and continuous

The Triple Whale Data Platform is an AI-optimized, fully-managed data warehouse that brings all of your data together as a single source of truth. This universal schema allows brands to conduct real-time data processing at scale.

2. Any platform can be connected via API

The ability to pull data from any major ad platform, 3PL and logistics applications, or any number of third-party integrations means all of your data is in one place and up-to-date.

3. Data is already organized in a way Moby AI can understand

No need to download or upload data at any time, it's all accessible with a simple prompt in Moby AI. An Agent can run from a single natural language prompt, and instantly produce a visual of your request.

4. Analyses can be run every day, all day

Unlike a standard LLM where you'd have to re-download, organize, then upload new data each day for a new analysis, Moby AI already has all of the information ready for the next query.

5. Moby AI knows ecommerce better than traditional LLMs

We already trained Moby AI on millions of ecommerce interactions and provided Moby AI the context to understand the specific needs, lingo, and nuances of ecommerce data. You don't have to spend time training the model to get the output you want.

6. You can create your own agent from scratch

If the 99+ pre-built ecommerce intelligence agents don't deliver your exact need, you can build an agent yourself using natural language prompts–quickly and easily.

7. Moby Agents can uncover anomalies a human might miss

Since agents can continuously monitor your performance, they are able to uncover anomalies in spend, performance, and site behavior far quicker than the average human, and can save you money in places you never would have found on your own.

8. Moby Agents take hours of manual work off of a team's plate

The estimated number of manual hours required for just the five examples above is 34-51 hours per month — essentially a full work week. For a brand running these analyses, that's 360-540 hours a year. When you consider how manual analysis can lead to delays in decision-making, the impact of lost time is even more clear.

Triple Whale

Final thoughts

While LLMs like ChatGPT and Claude can provide valuable insights when prompted correctly, the manual process of data collection, organization, and analysis remains time-intensive. For ecommerce brands that are serious about moving quickly to scale, Moby Agents can remove that complexity entirely to deliver instant, expert-level analysis without hours of manual labor.