



E B O O K

AI AS AN OPERATING LAYER.

From the experimental pilot to the implementation model: how to apply artificial intelligence across ecommerce operations.

STRATEGY • SOLUTION • DATA • AI • COMMERCE

AI IS NOT A PROJECT: IT IS A NEW OPERATING LAYER.

In recent years, many organizations approached artificial intelligence as an experimental initiative. Tests were run, tools were licensed, workshops were held, demos were built. **Yet a demo does not transform a company.**

Transformation happens when AI is embedded in everyday processes, helps make better decisions and shortens the distance between information and action. Within ecommerce, this opportunity is especially relevant: the digital operation produces an enormous amount of information.

Products

Content

Searches

Browsing

Transactions

Inventory

Customers

Campaigns

Inquiries

Reviews

Returns

Competitors

Promotions

Logistics

AI CAN INTERPRET, RECOMMEND AND AUTOMATE. BUT REAL VALUE TAKES FAR MORE THAN A TOOL: IT TAKES AN IMPLEMENTATION MODEL.


That model has three components that are rarely addressed together. The first is **the process**: which concrete task is being redesigned, who executes it today, how much time it consumes and what should change. The second is **the data**: which source feeds the solution, how reliable it is, who maintains it and what happens when information is incomplete. The third is **governance**: who validates, what can be automated without supervision, what requires approval and how everything is recorded.

When the first is missing, the organization buys technology without changing the way it works. When the second is missing, AI answers confidently about the wrong information. When the third is missing, risk appears: automations that amplify errors faster than any team can fix them manually.

This ebook walks that path in twelve chapters: first the four ways of applying AI and how to choose where to start; then concrete applications by area — catalog, content, service, data, marketing, forecasting, pricing and reporting; and finally governance, how to run a pilot and the maturity model that enables scale without losing control.

THE FOUR WAYS TO APPLY AI.

Talking about "using AI" without distinguishing formats creates the wrong expectations. A tool that helps write copy is not the same as a system that changes prices in production. Each format demands a different level of data, integration and supervision, and produces a different type of return. Ordering them lets you decide where to start and what to postpone.



01 ASSISTANTS

Help one person perform a task: write a description, summarize a report, analyze comments, propose keywords, prepare a presentation, organize information.

The human keeps the responsibility to execute and validate.

02 COPILOTS

Support recurring processes using specific context: a copilot for the commercial team, one for campaign analysis, another for querying metrics, for customer service or for catalog management.

03 INTELLIGENT AUTOMATIONS


Combine business rules with AI capabilities: classify and route tickets, detect anomalies and raise alerts, complete attributes, prioritize products, draft responses, identify causes of a drop.

04 AGENTS

Execute sequences of actions using systems and tools.

The level of autonomy must depend on the risk.

EXAMPLE · AN AGENT'S CYCLE




Most companies get their first results from the first two formats, because they require no deep integrations and are fast to measure. Agents, in contrast, only work when reliable data, documented processes and clear approval rules already exist. Starting there usually ends in an impressive demo that never reaches production.

HOW TO IDENTIFY REAL OPPORTUNITIES.

The first step is not asking "where can we use AI?". The right question is: **which processes are limiting the growth or efficiency of the business today?**

Identify processes with high volume, high frequency, manual work, recurring errors, repetitive decisions, unstructured information, dependence on a few people, long response times and commercial or operational impact.

PRIORITIZATION MATRIX



Document every use case with the same minimal format: which process is affected, who owns it, what information it needs, what is expected to improve, how it will be measured and what level of supervision it requires. That one-pager is what later allows comparing initiatives across areas without the decision depending on who argues best.

A frequent mistake is prioritizing by technological enthusiasm instead of impact. The most valuable cases are usually unspectacular: completing catalog attributes, classifying tickets, preparing weekly summaries, detecting anomalies. They make no headlines, but they free hours every week and are easy to measure.

AI APPLIED TO THE CATALOG.

The catalog is one of the areas with the highest automation potential. It is high-volume, highly repetitive work, with explicit per-channel rules and a direct impact on visibility, conversion and publication. The output is also easy to review: a person can validate in seconds whether a title complies or not.

That is why it is usually the best entry point. A team that takes weeks to prepare launch information can reduce that time measurably without touching core systems, and the learning it leaves behind — which rules must be made explicit, which sources are reliable, where review fails — is reused in every later use case.

TITLE GENERATION

Drafts based on product type, brand, attributes, benefits, keywords, length restrictions and channel rules.

DESCRIPTIONS

Adapts the same product for the own site, Amazon, Mercado Libre, Walmart, retailers, B2B, different markets and languages.

ATTRIBUTE ENRICHMENT

Identifies missing information from spec sheets, documents, images, descriptions, historical data and supplier information.

CLASSIFICATION

Suggests categories, subcategories, taxonomies, tags, collections and product relationships.

QUALITY CONTROL

Detects inconsistencies, duplicates, unauthorized claims, incomplete titles, contradictory information, wrong images and missing attributes.

AI can accelerate the work, but it must operate within brand, legal and commercial rules.

In practice, this means giving the model the same material you would give a new hire: brand manual, channel restrictions, allowed and forbidden claims, examples of approved content and examples of rejected content. Output quality depends far more on that context than on the model chosen.

CONTENT GENERATION AND OPTIMIZATION.

AI extends the content team's capacity. It should not be used only to produce *more pieces*, but to produce **better variants, with more speed and consistency**.

USE CASES

Descriptions	Category pages	FAQs	Buying guides
Comparisons	SEO content	Emails	Ads
Scripts	Marketplaces	Review responses	Campaign copy

PERSONALIZATION

The same product communicates differently by audience, funnel stage, category, need, channel, region, behavior and history.

HUMAN REVIEW IS MANDATORY

Claims · legal restrictions · medical information · sensitive technical data · commercial promises · brand communication · prices · promotional terms.

The simplest way to organize this is to classify content by risk. Low-risk content can be published with sample-based review; medium-risk requires channel-owner approval; high-risk always goes through specialized review. That criterion avoids two equally costly extremes: reviewing everything as if it were critical, or publishing without control.

CUSTOMER EXPERIENCE AND SUPPORT.

AI can act before, during and after an inquiry. Most implementations focus only on the "during" — the chat that answers — and that is why they disappoint: if product information is incomplete or the order status is unavailable, no assistant can resolve well.

The biggest return is usually at the extremes. Before, by reducing the number of inquiries generated. After, by understanding why they were generated and fixing the cause.

BEFORE

- FAQs
- Product information
- Guides
- Transactional messages
- Order status
- Delivery information
- Self-service

DURING

- Identify intent
- Classify the case
- Look up information
- Draft responses
- Translate and summarize
- Suggest next steps
- Escalate complex cases

AFTER

- Frequent causes
- Problematic products
- Sentiment analysis
- Group complaints
- Propose improvements
- Generate reports

THE GOAL IS NOT TO HIDE ACCESS TO A PERSON: IT IS TO RESOLVE SIMPLE CASES FASTER AND FREE THE TEAM FOR COMPLEX ONES.

Measure it with its own indicators: percentage of cases resolved without intervention, first-response time, total resolution time, reopen rate and satisfaction. If automatic resolution rises but reopens rise too, the system isn't resolving — it's postponing.

DATA ANALYSIS AND BUSINESS COPILOTS.

A copilot connected to business information changes how people access data. Instead of navigating multiple dashboards, a user could ask:

Why did sales drop yesterday?

Which categories are growing?

Which products are at risk of stockout?

Which campaigns are losing efficiency?

What are the main return reasons?

Which country is below target?

Which products need content?

Which promotions produced the most margin?

For the answers to be useful, the copilot needs reliable sources, consistent definitions, documented metrics, permissions, context, traceability and rules for acknowledging uncertainty.

THE CHATBOT DOES NOT REPLACE A DATA ARCHITECTURE. IT RELIES ON IT.

Before connecting a copilot, answer an uncomfortable question: how many different definitions of "sale", "margin" or "active customer" coexist in the company today? If there are several, the copilot won't unify them — it will expose them. Documenting metrics, agreeing on definitions and assigning owners per data domain is preparatory work that cannot be skipped.

Behavior under uncertainty also matters. A good copilot answers "I don't have that information" or "this data is incomplete" instead of filling the gap with a plausible estimate. The team's trust is built on that kind of answer, not on fluent text.

AI IN MARKETING AND GROWTH.

In marketing, the bottleneck is rarely a lack of ideas: it is the time between the idea and the measured execution. AI compresses that cycle on four fronts — segmentation, creative, optimization and experimentation — as long as an orderly campaign structure and a clear definition of "working" exist.

SEGMENTATION

Behavior · recency · frequency · value · categories · affinity · intent · churn risk.

CREATIVE

Create variations · adapt messages · propose concepts · switch formats · summarize benefits · localize content.

OPTIMIZATION

Audiences · campaigns · ads · landing pages · conversion · frequency · costs · profitability.

EXPERIMENTATION

Hypothesis generation · variant creation · documentation · analysis · prioritizing the next tests.

AI increases experimentation speed, but it does not replace a good strategy.

Multiplying variants without a hypothesis behind them produces noise: more pieces, more management overhead and no accumulated learning. The discipline that makes the difference is documenting what was expected to happen, what happened and what is decided as a result. AI can help with that too: drafting the experiment log and keeping it searchable.

FORECASTING, INVENTORY AND OPERATIONS.

AI can support planning by considering history, seasonality, promotions, trends, events, channels, prices, availability and external variables.

Demand forecasting

Stockout risk

Replenishment
recommendations

Anomaly detection

Order prioritization

Return prediction

Incident classification

Route optimization


Inventory recommendations

Models need monitoring. Consumer behavior changes, commercial conditions change and historical results do not guarantee the future. A forecast that was accurate for six months can degrade within weeks after an assortment change, a new competitor or an aggressive promotion.

That is why a forecast should not be delivered as a single number but with a range and a confidence level, and reviewed against reality every cycle. The value is not in always being right: it is in detecting early when the model stopped being right.

PRICING AND PROMOTIONS.

AI can analyze elasticity, competition, inventory, margin, demand, conversion, seasonality, events and commercial objectives. However, pricing is an area of **high impact and high risk**.



A good practical criterion: automation can execute on its own what an analyst would approve without hesitation and what can be reverted in minutes. Everything else — structural price changes, strategic products, categories with commercial agreements — goes through approval. And in every case, each change is logged with its reason, its owner and its subsequent result.

AUTOMATED REPORTING.

AI can turn information into narratives: weekly executive summaries, explanations of variances, risk and opportunity identification, comparison against targets, recommendations, presentation preparation and answers to questions.

FACTS

HYPOTHESES

RECOMMENDATIONS

INCOMPLETE
INFOCONFIDENCE
LEVEL

An automated report should never present an inference as a certainty.

AI GOVERNANCE.

Governance should not appear at the end: **it must be part of the design.**

EVERY SOLUTION NEEDS

An owner · an objective · authorized sources · defined users · permissions · usage rules · a validation process · metrics · monitoring · an escalation procedure.

FREQUENT RISKS

Incorrect information · inconsistent content · use of sensitive information · unsupervised decisions · bias · vendor dependence · lack of traceability · automations that amplify errors.

Governance should not be perceived as bureaucracy. Well designed, it is what enables speed: when the rules are clear, a team can launch a new use case in days, because it already knows which data it can use, who approves and how it is measured. Without those rules, every initiative gets debated from scratch.

HOW TO RUN A PILOT.

A pilot is not a test of the technology: it is a test of the redesigned process. That is why its design matters more than the tool chosen. Six steps are enough — as long as none is skipped, especially the baseline, the most-omitted step and the only one that later proves value.

01**Pick a specific problem**

Yes: "reduce the time needed to prepare product descriptions". No: "implement AI in marketing".

02**Define the baseline**

Measure current time, cost, volume, errors, quality and capacity.

03**Define the expected result**

Cut time by 50%, increase capacity, reduce errors, improve consistency.

04**Build the flow**

Input · processing · rules · review · approval · output · logging.

05**Measure**

Compare the result against the baseline.


06**Scale**

Only after proving value and stability.

A well-scoped pilot resolves in weeks, not quarters. If it can't be defined on one page, it probably isn't a pilot: it's a program, and it should be split.

AI MATURITY MODEL

No organization skips levels. The value of this model is not the label, but identifying which capability is missing to advance: usually accessible data from 2→3, governance from 3→4, and cross-area orchestration from 4→5.



C O N C L U S I O N

Artificial intelligence does not create value by existing. It creates value when it reduces friction, improves decisions and expands people's capacity.

THE QUESTION IS NOT HOW MANY TOOLS A COMPANY CAN ADOPT: IT IS WHICH PROCESSES IT CAN REDESIGN.

Organizations that treat AI as **an operating layer** will develop new capabilities. Those that treat it only as a trend will accumulate tests without transformation.



WWW.BALLOON-GROUP.COM · INFO@BALLOON-GROUP.COM