AI and Your Business

A Practical Guide for Business Leaders

A Note Before We Begin

This document is written for someone who runs a serious business, has no patience for jargon, and is rightly skeptical of claims that sound too good to be true. By the time you finish reading, you should be able to answer two questions for yourself.

First — what is AI and ML, really?

Second — where, specifically, could it move my P&L? Not in abstract terms. In terms of your actual business — the Amazon and Flipkart listings, the inventory decisions, the advertising spend, the returns headache — all of it.

There is no sales pitch here. Where the evidence is thin, this document says so. Where an idea sounds exciting but is genuinely hard to execute, this document says that too. The goal is a clear-eyed view of what is real, what is near, and what is still speculative.

One more thing. People often use "AI" as a catch-all term, as if it describes a single thing. It doesn't. The document uses **AI/ML** deliberately — Artificial Intelligence and Machine Learning — because these are related but distinct disciplines, and the distinction matters for how you think about applying them. More on that shortly.

Chapter 1

What AI/ML Actually Is – Starting With Something You Already Know

Algebra

Most explanations of AI start with terms like "neural network" or "transformer" — which mean nothing to someone who hasn't spent years in the field. So, let's start somewhere you've already been: Class 10 algebra.

Remember solving for unknowns? One unknown needs one equation. Two unknowns need two equations. Three unknowns, three equations. The logic is clean: as long as you have as many equations as unknowns, you can solve precisely.

Now consider this question: **what happens when you have a billion unknowns, but nowhere near enough equations to solve for all of them precisely?**

That is, more or less, the problem that machine learning is solving.

What the Unknowns Actually Are

Let's make this concrete with a real business example. Say you want to predict, for any given product listing on Amazon, how likely it is that a customer will return it.

Your inputs are everything you know: the product category, the price, the listing quality, the number of images, the review rating, the return rates for similar products, and dozens of other signals. Your desired output is a single probability: how likely is this item to be returned?

The "unknowns" are the weights — how much should each input factor matter? Is price more important than review rating? By how much? Does the number of images actually affect returns? How does a low price combined with a low review rating interact differently than either factor alone?

There are so many of these weights, and so many interactions between factors, that you cannot solve for them analytically the way you'd solve two simultaneous equations. There are too many unknowns and not enough neat equations.

The Brute Force Solution: Guess, Measure, Adjust, Repeat

Here is what machine learning does instead. It starts with a completely random guess for all the weights. Then it runs this guess against historical data — thousands of past orders where you already know what happened. It compares what the model predicted against what actually happened. The model will be terribly wrong at first. That's expected.

Then it makes tiny adjustments to the weights, in the direction that would have made the prediction slightly less wrong. It runs the data through again. Adjusts. Runs again. Over and over — sometimes hundreds of millions of times. Eventually, the weights converge on a combination that makes reasonably accurate predictions.

If you remember matrix multiplication — where you multiply two matrices to get a third — this is exactly the structure. You have Matrix A (your inputs: all the product and order data). You have Matrix C (your outputs: what actually happened). You are repeatedly tweaking Matrix B — the weights — until A × B produces something close to C.

> **Training is the process of finding the best Matrix B through iteration.**

> **The practical implication**
> Because you are finding the best approximation, the *model's output is always an educated guess, not a fact.*

Therefore, AI/ML implementations need thoughtful guardrails and scaffolding.

Why Google Knows More Than Your Local Startup

What separates a mediocre machine learning model from an excellent one is not primarily the algorithm. It is the number of attributes — the dimensions of information — that describe each thing being modeled.

Think of it like photography. A 1-megapixel image and a 40-megapixel image are both photos of the same subject. But the resolution is completely different. The 40-megapixel image captures details the 1-megapixel version simply cannot see.

A local e-commerce startup might build a customer recommendation model using 10 attributes about each customer: their age bracket, their city, their last purchase category, and a handful of other basics. Google and Meta are modeling each user across tens of thousands of attributes — every search query, every click, every scroll pattern, the time of day, the device, the content of ads they have lingered on, what they searched right before this. That is the difference between a 1-megapixel and a 40-megapixel model. Both are recommending products. The resolution is incomparable.

Deciding which attributes to include in a model is genuinely an art form. Two attributes that seem independent can actually cancel each other's effect out — a phenomenon statisticians call multicollinearity — and including both can make a model less accurate than including just one. Knowing which attributes are signal and which are noise, which combinations amplify each other and which cancel, is what data scientists spend their careers learning. It is not something a piece of software automatically figures out for you. This is why the quality of the team building the model matters as much as the tools they use.

The practical implication
The AI build-versus-buy decision is really a question about data. Build where you have information no one else has. Buy where someone else's data advantage is structural and permanent. Confusing the two is where most AI investment goes wrong.

The Needle in a Haystack — AI's Most Underrated Strength

One of AI/ML's most underappreciated capabilities is finding non-obvious solutions in enormous possibility spaces — what you might call finding the needle in a haystack.

Consider delivery route optimization. A city like Mumbai might have a driver making 80 stops. The number of possible sequences in which those 80 stops could be ordered is astronomically large — larger than the number of atoms in the observable universe. No human could evaluate more than a tiny fraction of those possibilities. Intuition gets you to a reasonable route. It cannot get you to the best route.

A machine learning optimization model can brute-force through vast combinations, evaluating millions of possible sequences and converging on solutions that humans would never find through intuition alone. A major logistics company implemented this kind of ML-based route optimization and reported saving tens of millions of miles driven per year. The efficiency gains did not come from humans being lazy — they came from the machine finding genuinely non-obvious combinations that no experienced route planner would think to try.

This pattern — brute-forcing through enormous possibility spaces to find non-intuitive optimal solutions — recurs across pricing optimization, inventory allocation, advertising bid management, and many other business problems. It is one of the clearest areas where machines have an inherent advantage over humans, regardless of experience.

Chapter 2

A Map of the Territory — The Different Kinds of AI/ML

One of the most common mistakes business leaders make is treating AI as a single thing. It is not. Saying "we should use AI" is a bit like saying "we should use medicine" — technically true, but not a decision. There are dozens of distinct types of AI/ML systems, each suited to different problems, requiring different data, and delivering different kinds of outputs.

Here is a practical map of the landscape. The goal is not to make you a technical expert — it is to give you enough vocabulary to have informed conversations with people who are building these systems for you.

Type 1: Classical Machine Learning — The Workhorse

This is the oldest and most battle-tested category, and in many ways the most practically reliable. Classical ML takes structured data — rows and columns, like a spreadsheet — and learns to make predictions from it.

Examples include demand forecasting (predicting how much of a product you will sell next month), fraud detection (predicting which transactions are likely fraudulent), customer churn prediction (predicting which customers are likely to stop buying), and dynamic pricing models (predicting how demand will respond to price changes).

Google and Facebook have been running billion-dollar businesses on classical ML for over fifteen years. It is not glamorous, it does not write sentences, and it will never go viral on LinkedIn — but it consistently generates measurable financial returns when implemented on clean data with clear objectives. This is the category where the proof of concept is most mature.

Type 2: Computer Vision — Teaching Machines to See

Computer vision models are trained to interpret images and video. They can identify objects, assess quality, count inventory, read text from photos, and increasingly, generate or enhance images.

Type 3: Large Language Models — What Everyone Is Talking About

Large Language Models, or LLMs, are the systems behind ChatGPT, Claude, Gemini, and similar products. They are trained on an extraordinary volume of text — trillions of words from books, articles, websites, research papers, and more. Through this training, they develop the ability to read, write, summarize, translate, and reason through language.

This is genuinely new capability. Earlier ML systems were narrow: a fraud detection model could only detect fraud. An LLM can help write a product description, then summarize a contract, then draft an email to a supplier, then answer a customer question — all in sequence. It is a general-purpose language tool rather than a specialized predictor.

But — and this is important — general-purpose does not mean infallible. LLMs have a specific failure mode that every business leader deploying them needs to understand clearly.

How LLMs Actually Work — And Why They Sometimes Say Wrong Things With Complete Confidence

When an LLM generates a sentence, it does not retrieve that sentence from a database. It generates it word by word, with each word being a probabilistic choice based on everything the model has learned.

Here is the clearest way to think about it. An LLM trained on trillions of sentences has encountered the phrase "The Bombay Stock Exchange is located in..." tens of thousands of times, and in the overwhelming majority of those instances, the next word was "Mumbai." But it has also seen older texts — from before 1995, when the city was officially renamed — where the same sentence continues with "Bombay." In the model's learned distribution, "Mumbai" might appear with 99.99% probability and "Bombay" with 0.01% probability.

When the model generates a response, it samples from this distribution. Almost always, it says Mumbai. Very occasionally — rarely, but not never — it might say Bombay. That is a relatively harmless example of what we call a **hallucination**: an output that is confidently stated but factually incorrect, not because the model is broken, but because it is sampling from a distribution rather than looking up a verified fact.

Now here is where it becomes genuinely dangerous. An LLM has also seen hundreds of thousands of legal documents, and it has learned exactly how a legal citation looks: "In the case of [Name] v. [Name] ([Year]), the court held that..." It knows this structure perfectly. If asked to write a legal brief, it will generate things that look exactly like properly formatted case citations — with realistic names, realistic years, realistic docket numbers — whether or not those cases actually exist.

In 2023, two lawyers in the United States filed a court brief citing six specific cases as legal precedents. Every single one was fabricated. The cases did not exist. The lawyers had used an LLM to help draft the brief and had not checked the citations. The court discovered the fabrications. The lawyers were sanctioned. One faced disbarment proceedings. The story made international news.

This is not a rare malfunction. It is a predictable consequence of using a probabilistic text generator for a task that requires verified facts. The model did exactly what it was designed to do — generate fluent, legally-formatted text. The problem was a category error: the wrong tool was used for the wrong job, without the appropriate human verification layer.

> **The practical implication**
> AI implementation is a design problem that spans technology and *process*.

When Randomness Is a Feature, and When It Is a Bug

The probabilistic nature of LLMs is genuinely valuable in certain contexts. When you ask a model to brainstorm twenty product names, you want it to range across the probability space and surface unexpected combinations. When you ask it to draft five versions of a product description or marketing copy for A/B testing, you want variation. When you ask it to identify themes in ten thousand customer reviews, you want it to surface patterns you hadn't thought to look for.

In these contexts — creative generation, exploratory analysis, summarization, ideation — the randomness is the point. You are not asking for facts; you are asking for possibilities.

But when a task requires verified, specific facts — legal citations, financial figures, regulatory provisions, technical specifications — you need a system that retrieves verified information from a trusted source, not one that generates plausible-sounding text from a probability distribution. The architecture needs to be different, and human review is non-negotiable.

The Most Important Mental Model for Using AI/ML in Business
AI does not replace human judgment. It elevates what human judgment is spent on. The businesses getting the most consistent value from AI are the ones that have redesigned their workflows around a simple shift: humans stop being producers and become editors and supervisors. The machine handles volume, pattern, and execution. The human provides taste, strategy, and accountability. The right question is not "how do I automate this?" — it is "where does a human need to be in the loop, and for what purpose?"

AI/ML as a Complement — Not a Replacement

The point that gets lost in the headlines about AI replacing jobs is that the most consistent and well-documented financial impact of AI/ML in business operations has come not from replacing people, but from making existing people significantly more capable.

A customer service agent with an AI copilot handles more queries per hour, produces more consistent responses, and can serve customers she'd never be able to reach alone. A buyer with an ML-powered demand forecasting tool makes better inventory decisions with less effort and time. A content writer with an LLM drafts ten product descriptions in the time it used to take to write two. A logistics planner with route optimization software covers more stops with fewer resources.

In each case, the human is still doing the job. The AI is doing the rote, the repetitive, the pattern-matching — freeing the human to focus on judgment, relationships, exceptions, and the things that actually require a person.

This is the framing that tends to resonate most strongly with management teams: not "AI will replace your people" but **"AI will allow your existing team to do the work of a team twice its size."** Senior people spend time on senior work. Junior people are backed by a system that catches errors and surfaces relevant information. New hires get up to speed faster because the AI acts as an institutional memory they can query.

The goal is not fewer people. It is more leverage from the people you have.

ChatGPT Is Not AI. AI Is Not ChatGPT.

The final thing worth establishing before we move to business applications: ChatGPT — or Claude, or Gemini, or any specific product you have heard of — is one narrow application of one type of AI system. It is the equivalent of saying "medicine" when you mean "a specific brand of ibuprofen."

The AI/ML landscape is a discipline with dozens of branches, each applicable to different problems, requiring different data, different expertise, and different organizational practices. A well-designed AI/ML program for a business like yours will likely use multiple types of systems for different purposes: classical ML for forecasting, computer vision for imagery, LLMs for content and customer service. They are not interchangeable.

The practical consequence of this is that there is no single AI product you can purchase and deploy that solves all problems. The discipline requires mapping specific problems to appropriate techniques – and combinations of techniques — and that mapping is where most of the expertise lies.

Chapter 3

Where AI Is Already Working – Real Evidence From Real Businesses

The examples below are real, documented outcomes from businesses operating in competitive markets — not laboratory experiments or vendor marketing claims.

Customer Service: The 14% Productivity Gain

A large-scale field study of customer service agents using an AI assistant found that agents handled approximately 14% more customer issues per hour when the AI was actively helping them[1]. Crucially, the gains were not uniform: less experienced agents saw improvements of up to 34%, while the most experienced agents saw smaller gains. The AI was effectively transferring the knowledge and response patterns of the best agents across the entire team.

This pattern — AI functioning as a knowledge transfer mechanism that raises the floor across a team — recurs in most documented customer service AI deployments.

A Live Example: Marvin by Aegion
Our own customer support agent, Marvin, is currently handling approximately 60% of inbound customer support requests autonomously for our clients — answering queries, retrieving information from knowledge bases, and resolving routine issues without any human involvement. For the remaining 40% that require human judgment, Marvin packages the full context of the issue before escalating: the customer's history, the relevant documentation, the prior conversation, and a summary of the issue. The human agent receives a complete briefing

and can focus entirely on making the decision — not on reconstructing what happened.

Returns: Size Guidance With AI

Zalando, Europe's largest online fashion retailer, built an in-house AI system that combines purchase history, customer return reasons, brand-provided measurements, and fit feedback from a dedicated team who physically try on items. The system flags sizing anomalies — "runs small, size up" — and provides personalized size recommendations based on each customer's history. In 2023, they extended this further, allowing customers to generate body measurements from two smartphone photos.

The result: size-related returns dropped[2] by 10% for items where size advice was offered, compared to the same categories without it. Notably, this is not a year-on-year comparison but a controlled within-category comparison — the more rigorous test.

The financial significance is easy to understate. Roughly one third of Zalando's total returns are size-related, and at a platform processing tens of millions of orders annually, a 10% reduction in that category represents a material shift in reverse logistics costs. The fix was not changing the products. It was changing the information available to customers before they bought.

Route Optimization: The Non-Obvious Solution

UPS, a major global logistics company, implemented ML-based route optimization across its delivery network and reported saving hundreds of millions of miles driven per year. These gains did not come from drivers discovering something they had missed. They came from the algorithm finding route combinations that were genuinely non-intuitive — sequences that no experienced dispatcher would have generated by hand, but which consistently outperformed manually planned routes. This is the needle-in-a-haystack capability described in Chapter 1 at work: brute-forcing through enormous possibility spaces to find optimal solutions humans cannot reach through experience alone.

Writing and Content: 40% Time Reduction

A randomized experiment published in *Science* assigned professional writing tasks — press releases, persuasive essays, detailed analyses — to 453 college-educated workers, with half given access to AI assistance. The results were striking: time to completion dropped by 40% on average and independently evaluated output quality rose by 18%[3]. Crucially, the gains were largest for lower-performing writers. The AI effectively compressed the gap between the best and average communicators.

For a business producing product descriptions, promotional copy, or customer communications at scale, this has a concrete implication: the same team produces more, faster, and at more consistent quality. The bottleneck in content production shifts from writing to editing and judgment — which is exactly where human time is better spent.

The Emerging Frontier: AI Agents

Everything described so far — forecasting, content generation, image analysis, sentiment classification — involves AI producing an output that a human then acts on. A forecast that informs a purchase order. A draft that a writer edits. An alert that a manager reviews.

> **AI agents are different. They do not just produce outputs. They take actions.**

An agent is an AI system that can observe its environment, reason about what it sees, take an action in the world — and then observe the result of that action and decide what to do next. It operates in a loop, not in a single pass. This is what distinguishes an agent from a model: a model answers a question; an agent pursues a goal across multiple steps, using whatever tools it has access to.

To make this concrete: the difference between an AI that drafts a response to a customer complaint and an AI agent that handles the complaint is the difference between a word processor and a member of your customer service team. The word processor produces text. The agent reads the message, looks up the customer's order history, checks the product's known issue database, decides whether this qualifies for a proactive replacement, drafts a personalized response, sends it through the correct channel, and updates the ticket — without a human touching any of it.

How an Agent Works

At its core, an agent runs a loop that has three stages, repeated as many times as needed:

- **Observe:** The agent reads its inputs. These might be a new customer message, a product review that just appeared, a change in a competitor's price, or the result of a database query it ran in the previous step.
- **Reason:** The agent decides what to do next. Should it respond immediately, or look up more information first? Is this query within its competence, or does it need to escalate to a human? This reasoning step is where the underlying language model earns its keep — it is what separates an agent from a simple rule-based automation.
- **Act:** The agent takes an action using one of its available tools. Searching a knowledge base. Querying an order management system. Sending a message. Creating a

ticket. Updating a record. Each action changes the state of the world in some small way.

The loop continues until the agent either completes the task, reaches the limits of its competence and escalates, or decides no action is needed.

The tools an agent can use are what give it reach. A customer service agent with access to your order management system, your returns database, your product knowledge base, your messaging platform, and your escalation channel can handle a far wider range of situations than one with access to only a scripted FAQ. Integrations are not just a convenience feature in agentic systems — they are what determines the scope of what the agent can actually do.

The Autonomy Spectrum

Not all agents operate with the same degree of independence, and where you position an agent on the autonomy spectrum should be a deliberate choice based on the stakes involved.

At one end, a fully supervised agent can propose actions for human review prior to execution. The agent does most of the work, but a human controls every output.

Further along the spectrum, an agent operates autonomously within defined boundaries — handling routine cases independently but escalating anything that falls outside those boundaries. It might respond to the eighty percent of queries that fit known patterns and flag the remaining twenty for human attention. This is where most production agentic systems eventually settle.

At the far end, a fully autonomous agent makes and executes decisions independently, within whatever guardrails have been built in. This is appropriate for some tasks — a repricing agent that adjusts prices within margin guardrails does not need human approval for every adjustment — but requires significantly more investment in the guardrails themselves before it should be trusted.

The right place on this spectrum is not "as autonomous as possible." It is whatever level of autonomy the task's stakes and your integration quality actually justify.

Why Agents Are Both More Powerful and More Complex

The same properties that make agents powerful also make them harder to implement well.

A standard AI model produces an output. If the output is wrong, a human catches it and corrects it. The error is contained. An agent takes actions. If an early action in a multi-step sequence is wrong, subsequent actions are built on a flawed foundation. The error compounds. An agent that misidentifies a customer's issue and then sends a proactive outreach message based on that misidentification has not just made one mistake — it has taken an incorrect action in the world that may be difficult to reverse.

This means the implementation bar for agentic systems is meaningfully higher than for analytical AI. Getting the data integrations right matters more, because the agent is acting on what it retrieves. Getting the escalation logic right matters more, because the agent needs to know the boundaries of its competence. Getting the access controls right matters more, because an agent with too much access and too little oversight can do real damage at the speed of software.

The reward for getting these things right, however, is proportionally larger. A well-implemented agent can handle entire workflows end-to-end, at a scale and consistency that no team of humans could match. It does not get tired, does not forget context from earlier in a conversation, and can run across hundreds of simultaneous interactions without degradation. This is why customer service — high volume, high repetition, context-dependent, with clear escalation triggers — is where agentic systems are delivering the clearest early returns.

> **What This Means for Chapter 4**
> Several of the applications described in the next chapter are agentic systems in their full form, not just analytical tools. These are flagged where relevant. They are consistently in the higher implementation complexity categories — not because the underlying AI is more sophisticated, but because the integration depth, escalation logic, and access controls required to deploy them responsibly are substantially more demanding than a model that generates a report for a human to review.

The Commoditization Warning

There is a broader trend worth flagging here, particularly for marketplace sellers. The platforms themselves — Amazon, DoorDash, and others — are using the same AI/ML tools that sellers are using. They are deploying ML to identify their best-performing products, to optimize search rankings, to surface competitive alternatives to buyers, and increasingly to inform their own private-label strategies.

This means that over time, competing purely on product and price on a marketplace becomes increasingly difficult. The platform has too much information, too many levers, and too much incentive to commoditize the seller relationship.

The businesses that will be most resilient over the next five years are the ones building a direct relationship with their customers — a brand voice, a content presence, and traffic channels that they own. AI/ML makes this more accessible than ever: creating personalized, high-quality content at scale is no longer the exclusive domain of large marketing teams. A business of any size can now produce content output — blog posts, short-form video scripts, product guides, social media — that would previously have required a team of ten.

We will return to this in the P&L section with specific tactical ideas. For now, the principle: every marketplace seller should be thinking about their owned channels alongside their marketplace presence, and AI/ML makes building those channels substantially more affordable.

Chapter 4

Where AI/ML Moves the Numbers — A Domain-by-Domain View

This chapter maps the AI/ML opportunity across eight operational domains. Each section explains what the domain is, illustrates one specific way AI creates impact, and identifies where implementations typically go wrong. Difficulty ratings reflect the realistic entry point for each domain — the first thing you would actually do, not the most ambitious thing you could eventually build.

Difficulty ratings are indicated as follows:

- 🟢 *Low Hanging Fruit — can be implemented quickly with available tools*
- 🟡 *Moderate — requires some data engineering*
- 🔴 *Involved — significant build, needs dedicated technical resources*

4.1 Catalog and Listing Health

Your catalog is the foundation everything else rests on. A suppressed listing generates zero revenue regardless of how strong your advertising or inventory position is — and suppression is rarely announced loudly. Content gaps, missing compliance attributes, and listing quality failures erode availability rates, search ranking, and conversion silently, often for weeks before anyone notices.

AI makes catalog health a continuous process rather than a periodic campaign. LLMs audit every listing against current platform guidelines, identify missing mandatory attributes, rewrite titles and bullets for search term alignment, and generate A+ content at scale.

Computer vision models produce contextual lifestyle images — the pressure cooker in an Indian kitchen, the storage unit in a styled bedroom — at a fraction of traditional photography costs, and at a quality level that consistently improves conversion over white-background shots alone. Across a catalog of thousands of SKUs, the difference between AI-assisted catalog management and manual review is the difference between comprehensive coverage and a perpetually incomplete backlog.

Example: A catalog of 4,000 SKUs might contain 200 suppressed listings, 800 with incomplete attributes reducing search visibility, and 150 with non-compliant images — all invisible to the team until sales velocity drops. An AI audit produces this picture in hours, ranked by revenue impact of the affected SKUs. In one case, a team discovered that a missing regulatory attribute across 300 FMCG listings had quietly taken them inactive for weeks, recoverable in a single correction cycle once identified.

The implementation risk is subtle rather than dramatic: AI-generated content that sounds authoritative but contains an incorrect specification — a wrong capacity, a mislabeled compatibility claim — reaches customers and generates returns and negative reviews before anyone notices. The model must be grounded in verified product data, not given latitude to generate creatively. Human review of AI-generated content before publishing is not a courtesy; it is the control.

🟢 **Low Hanging Fruit**

4.2 Promotional Strategy and Price Integrity

Promotions and pricing are two sides of the same margin equation and managing them in isolation is how margin quietly disappears. Configuring deal slots on the wrong SKUs, at the wrong discount depth, without tracking whether price integrity holds across channels and time — this is a continuous optimization problem that exceeds the capacity of manual monitoring at any meaningful catalog scale.

AI addresses both sides simultaneously. Promotional performance models analyze the historical return of every deal mechanic by category and discount depth, distinguishing between configurations that generate incremental margin-positive volume and those that pull forward demand, spike return rates, or simply subsidize customers who would have bought anyway.

On pricing, dynamic repricing models monitor competitor positions and market signals in real time, adjusting within guardrails you define — never below a margin floor, never more than a defined spread above category median — maintaining Buy Box competitiveness without requiring manual intervention.

Example: Historical analysis of a deals program in an electronics category shows volume spikes at three times normal ACoS, accompanied by a measurably higher post-

promotion return rate. On a full P&L basis, the promotions are margin-negative. The same budget applied to coupons on FMCG SKUs generates a 2x conversion lift with no incremental return cost. The reallocation is obvious in hindsight — and invisible without the analysis.

Where implementations go wrong: a repricing model without carefully defined margin floors will chase competitors into unprofitable territory, and it will do so at a speed and scale that makes the damage hard to detect until it has already accumulated. The guardrails require at least as much attention as the model itself — and they need regular review as cost structures change.

Moderate – Requires Some Data Engineering

4.3 Inventory Health and Availability

Inventory is the largest capital commitment in most businesses, and the cost of getting it wrong runs in both directions simultaneously. Stockouts lose sales. Excess inventory accrues storage costs, forces margin-dilutive markdowns and ties up working capital that could be deployed elsewhere. Both errors look identical from the outside — as missed profit — and both are preventable with better information.

ML-based forecasting models learn the patterns that drive demand at the SKU level: festival calendars, regional variation, day-of-week effects, promotional lifts, and competitor activity. The same data infrastructure supports a continuous early warning system that monitors inventory trajectories in real time, flagging stockout risk and aging inventory weeks in advance. The goal is not a better spreadsheet — it is the shift from reacting to problems after they occur to intervening before they do.

Example: A cleaning product variant has a consistent demand spike in the second week of April tied to spring cleaning, running 40% higher in Northern states than elsewhere. An ML model identifies this pattern across multiple years of history, recommends pre-positioning inventory in the relevant fulfilment centers three weeks ahead, and simultaneously flags a different category where current sales velocity puts the business on pace to incur long-term storage fees in twelve days — with a specific coupon depth recommendation to clear the inventory in time.

The implementation challenge is almost always upstream of the model. Inventory forecasting requires clean, connected data from systems that were not designed to talk to each other — sales history, promotional calendars, replenishment lead times, and current stock levels across multiple warehouses, each with different update frequencies and identifiers. Reconciling these into a coherent, reliable data feed is where most implementations stall, and it requires more time and care than the forecasting model itself.

Moderate – Requires Some Data Engineering

4.4 Supply Chain and Vendor Performance

Upstream supplier decisions determine downstream business outcomes in ways that are rarely visible until something goes wrong. A supplier whose lead times have quietly lengthened, a batch with a defect pattern not yet visible in aggregate return data, a vendor whose pricing has drifted above market — these are the slow sources of margin erosion and service failures that do not appear in weekly reports until they have already compounded.

AI brings systematic visibility to supplier performance by integrating data that currently lives in separate places: purchase orders and receipt timestamps for lead time tracking, quality inspection logs and return reason codes for defect attribution by supplier and batch, and review sentiment analysis that captures product quality signals before they work their way through the formal QA process. The output is a continuously updated performance picture for each supplier — factual, documented, and directly useful in procurement conversations.

Example: Review analysis for an apparel SKU shows a rising volume of stitching complaints over a six-week window — a pattern invisible in the aggregate return rate, which lags by several weeks. AI traces the complaints to a specific supplier batch using lot codes from the returns management system. The team approaches the supplier with documented evidence — not a vague quality concern, but a correlation between batch, complaint type, and return rate — and resolves the issue before it compounds further into review score damage.

The risk in vendor analytics is attribution error: incorrectly assigning blame to a supplier for a problem that originated in logistics handling, listing inaccuracy, or a seasonal effect. Models that do not account for confounding variables can produce supplier scorecards that are precise but wrong. Acting on them damages supplier relationships that can take months to repair and diverts attention from the actual root cause.

Moderate – Requires Some Data Engineering

4.5 Customer Acquisition and Advertising

Lower funnel advertising is a high-frequency optimization problem at a scale that exceeds human capacity to manage manually. Running multiple campaigns across brands, products, geographies, and channels generates more optimization decisions per day than any team can meaningfully process. The consequence is systematic: overspending on non-converting search terms, underbidding during high-conversion windows, and leaving organic content — the traffic that compounds rather than expires — largely undeveloped.

AI applies across the full acquisition stack. In paid advertising, models continuously analyze search term performance, prune non-converting spend, and adjust bids by time and conversion probability — effectively running the optimization work of a large team on a continuous basis. In organic content, LLMs generate SEO-optimized articles, buying guides, and product video scripts at a scale that builds owned traffic permanently. A business that previously produced four Reels per month, bottlenecked by scripting time, can produce twenty with the same creative resources when AI handles the scripting. A blog post that ranks for a relevant search query generates traffic indefinitely without ongoing spend.

A broad match campaign on "running shoes" accumulates spend against "running shoe reviews," "how to clean running shoes," and "running shoe size guide" — informational queries from people with no purchase intent. Meanwhile, "men's lightweight trail running shoes size 11" is converting well in broad match but competing against unrelated terms in the same budget pool. AI identifies the wasteful terms and adds them as negatives, then surfaces the high-converting long-tail and promotes it to an exact match campaign with a dedicated bid. ACoS drops without losing meaningful volume. The same analysis runs continuously across every campaign in the account.

The failure mode in advertising automation is optimizing for the wrong metric. A model that minimizes ACoS without accounting for lifetime value, margin by category, or post-promotion return rates can be technically correct and commercially damaging simultaneously. The model optimizes for whatever objective you give it with precision — which means the objective requires more business judgment to define than the model requires technical skill to build.

🟡	**Moderate – Requires Some Data Engineering**

4.6 Returns, Customer Service, and Reputation

The post-purchase experience is where margin is made or destroyed in ways that rarely appear clearly in standard reporting. Returns carry a true cost distributed across reverse logistics, processing, refurbishment loss, and review score impact — typically 30 to 50% of product value for items that cannot be resold as new. US retail returns alone were projected at $890 billion in 2024, with 16.9% of sales returned across the industry. Customer service volume that should be preventable consumes team capacity. Review scores, once damaged by a systematic product or experience issue, are slow and expensive to recover.

AI works on this cluster of problems simultaneously. Return reason analysis identifies the specific attributes — size chart accuracy, packaging integrity, content gaps — driving recoverable returns, creating a feedback loop back to catalog and supplier decisions. Aspect-level sentiment analysis on reviews surfaces the product dimensions dragging ratings down — not the aggregate score, which tells you nothing actionable, but the specific attribute receiving 2 stars while others receive 4. And customer service automation handles the 60 to 70% of inbound queries that are genuinely repetitive, freeing the team for the cases that actually require judgment. In a documented field study, agents using AI assistance handled approximately 14% more queries per hour, with the largest improvements among less experienced team members — the AI effectively transferring the knowledge of your best people across the entire team.

Example: A formal shirt SKU carries a 22% return rate. Aggregate data says 60% are 'did not fit.' AI clusters return comments and identifies three distinct sub-issues: a size running small, a collar width labelled incorrectly, and model photography that misrepresents the cut. Each fix is different and none requires changing the product. The listing is updated, the size chart corrected, and the supplier briefed on the collar specification. Return rate drops to 14% over two months — without a single product change.

The implementation complexity in this domain spans a wide range. Review mining and return analysis are accessible starting points that can generate insight quickly. Customer service automation at the level of reliability customers expect — where a wrong answer damages a relationship rather than being a neutral non-event — requires solid integrations across order management, product knowledge, and escalation logic. Starting in supervised mode, where the system drafts and a human approves before anything reaches the customer, is the right first step before expanding to autonomous operation.

🟡	**Moderate – Requires Some Data Engineering**

4.7 Business Intelligence and Reporting

Most businesses at meaningful scale are making decisions based on data that is days old, partially reconciled, and aggregated in ways that obscure what matters. The SKU-level P&L — accounting for commissions, fulfillment fees, advertising spend, return costs, and COGS in a single view — exists in theory but requires hours of manual work to produce. It happens infrequently, so decisions get made without it. Products that appear profitable in revenue terms are losing money after fees. Products that appear minor are disproportionately valuable. These are not edge cases — they are the norm in businesses where the data infrastructure has not kept pace with the catalog.

AI automates the reconciliation and reporting layer, connecting data from multiple sources into a daily view that updates without manual intervention. The same infrastructure supports fee anomaly detection — systematically identifying where platform charges deviate from expected values — and forward-looking cash flow forecasting that incorporates demand forecasts, procurement schedules, and platform disbursement timelines. The goal is to make the information that currently requires a finance analyst two days per month available every morning, to anyone who needs it.

Example: Amazon calculates FBA fees from dimensional weight measurements recorded in the fulfilment center. If those measurements are wrong — which happens, particularly after packaging changes — you are being systematically overcharged on every order for that SKU. Manual detection requires someone to notice an anomaly in per-unit economics across a large and noisy dataset. AI detects the pattern across thousands of orders, quantifies the total overcharge, and generates the documentation needed to file a dispute — recovering fees that would otherwise never be recovered.

The implementation caveat here is organizational rather than technical. Automated reporting is only useful if the decisions made from it are disciplined. A daily SKU-level P&L that is reviewed monthly has not solved the problem — it has just made the problem more visible without changing the outcome. The tools are genuinely accessible; building the decision cadence and accountability structure around them is the harder and more consequential work.

Low Hanging Fruit

4.8 Business Process Automation — The SG&A Opportunity

Every section above addresses how AI affects the commercial operations of a business — the revenue lines, the cost of goods, the customer experience. These are the obvious targets, and the conversation usually stops there. But there is an equally significant set of opportunities that most AI discussions never reach: the SG&A layer.

For a business operating at meaningful scale, SG&A — finance operations, HR, legal and compliance, IT support, marketing operations, internal reporting — typically represents 15 to 25% of revenue. This is the overhead of running the business itself. It is also where the next wave of AI impact is being most consistently underestimated.

From BPO to BPA

In the 1990s and 2000s, Infosys and others developed an important insight: large organizations were paying expensive people to do repetitive, rule-based back-office work. Business Process Outsourcing was the answer — move the work somewhere cheaper. The pitch was pure labor arbitrage, and it worked for two decades.

Business Process Automation is the next evolution of the same insight. BPO asked: who is the cheapest person to do this task? BPA asks: does this need to be a human task at all? A vision model reads the PDF invoice, extracts structured data, matches it against the purchase order, codes it to the correct GL account, and flags only the exceptions that require human review — in seconds, at any volume, with a complete audit trail. The logic is the same as BPO — remove cost and friction from back-office work — but the mechanism has advanced from labor arbitrage to automation. Every generation of technology has pushed repetitive work further down the cost curve. BPA is the current step in that progression.

The Scale Asymmetry
A human accounts payable team of five processes roughly 500 to 800 invoices per month at acceptable quality. An AI system handles 50,000 with the same accuracy and flags the same exceptions for human review. The team of five still exists — they manage exceptions and vendor relationships. They stop being the bottleneck.

The Five Domains

In finance and accounting, the highest-volume opportunities are invoice processing and three-way matching, expense report compliance, financial report drafting from structured data, and payment fraud detection. In human resources: candidate screening and CV analysis, onboarding and HR policy Q&A, and performance review drafting. In legal and compliance: contract clause extraction and first-pass redlining, and regulatory change monitoring across categories and geographies. In IT and internal operations: helpdesk automation for repetitive ticket categories, and knowledge management that makes institutional knowledge queryable rather than siloed in individuals. In sales and marketing: CRM data hygiene, proposal and RFP generation, and lead scoring from historical conversion patterns.

Moderate – Requires Some Data Engineering

Chapter 5

The One Investment That Determines Whether Any of This Works

Everything described in Chapter 4 is theoretically available to you. The tools exist. The techniques are proven. The question is why so many businesses attempt AI/ML programs and see disappointing results.

The answer, almost universally, is not the model. It is the data. Studies of enterprise AI implementations consistently find that **between 60 and 85% of AI projects fail to reach production or fail to deliver expected value** — and the most frequently cited reason is not algorithmic complexity. It is data that is inconsistent, inaccessible, fragmented across systems, or simply not organized in a way that a model can use.

Why Data Is Harder Than the Model

Building a machine learning model, given clean and well-organized data, is increasingly a commodity activity. The tools are excellent. The techniques are well-documented. A capable data scientist can have a working prototype running in days. Cleaning and organizing the data to make that model possible — that is the hard, expensive, time-consuming part. And it is almost always harder than the business expected.

Here is a concrete illustration of what this looks like in practice. Say you want to build a model that predicts which customers are likely to churn. You need to combine your CRM data (account history and relationship notes), your billing system (payment patterns and contract terms), your support platform (ticket volume and sentiment), and your product or service data (usage and engagement). Each of these systems uses different identifiers for the same customer. Dates are formatted differently. Some systems record activity at the account level, others at the individual contact level. What one system calls "at risk" another calls "escalated." What the support tool flags as "resolved" may still be open in the CRM.

None of this is insurmountable. But none of it is trivial either. And until it is resolved, your churn prediction model is running on inputs that may look like data but are actually a patchwork of inconsistencies — producing outputs that are confidently wrong in ways that are difficult to detect until the damage is done.

What This Means Practically

The most durable investments you can make in preparation for AI/ML are not tools. They are infrastructure — and that infrastructure has two distinct components that are easy to underestimate until you are mid-implementation.

The first is data. Your customer records should use consistent identifiers across every system that touches them — your CRM, your billing platform, your support tool. Your operational data should capture outcomes in structured fields, not free-text comments that no model can reliably parse. Data from different systems — sales, marketing, finance, operations — should be linkable at the transaction level, not just in aggregate weekly reports. These are not glamorous investments. They are also not optional. The businesses that made them three years ago are deploying AI/ML in weeks. The businesses that did not are still explaining to their data scientists why the model outputs do not match reality.

The second is process. Clean data is necessary but not sufficient. AI generates outputs — forecasts, draft responses, flagged anomalies, suggested actions — and someone in your organization needs to know what to do with them. Which outputs get reviewed by whom, at what frequency, with what authority to act? Where does a human need to be in the loop, and where can the system run autonomously? These workflow questions need to be designed deliberately. Organizations that deploy AI without answering them find that the tool sits underused, or worse, that its outputs are ignored because nobody owns the responsibility of acting on them.

This is where most AI/ML programs quietly stall — not because the model failed, but because the surrounding infrastructure was never built. The technology is the easiest part. The data engineering and process design that make it useful in a real business are where the actual work lies — and where the difference between a successful implementation and an expensive pilot is made.

> **The Key Point**
> Models come and go. Data infrastructure, once built, is an asset that compounds. The business that has clean, accessible, well-organized data today will be able to adopt whatever the next generation of AI/ML tools looks like, because the foundation is already in place. The business that does not will face the same data cleaning problem with every new implementation — paying for it repeatedly, in time and money.

Chapter 6

What Can Go Wrong — Risks and How to Manage Them

Every technology that creates meaningful upside also creates meaningful downside if implemented carelessly. AI/ML is no exception. In fact, because AI/ML operates at scale and often with some degree of autonomy, mistakes compound faster than they would in a manually operated process.

This chapter is not intended to discourage you from pursuing AI/ML. It is intended to ensure that when you do, you are doing so with a clear-eyed view of what can go wrong — and with the right safeguards in place from the beginning.

These are the risks that matter most for a business like yours.

Risk 1: Your Data Leaving Your Business

When employees use third-party AI tools — ChatGPT, Copilot, Gemini — and paste in customer records, supplier contracts, pricing strategies, or internal analyses, that data is transmitted to an external provider's servers. Depending on the tool, the account type, and the terms of service, it may be used to train future models. This is not hypothetical: several enterprise data exposures have been traced directly to employees using consumer AI accounts rather than enterprise-licensed versions with appropriate data agreements.

The consequence is twofold. Commercially sensitive information — your negotiated supplier terms, your margin structure, your customer contracts — may leave your control permanently. And if personal data is involved, you may have a regulatory exposure under DPDPA or GDPR regardless of whether any harm results.

Mitigations: Establish a clear policy on which data categories may and may not be used with external AI tools. Ensure any AI vendor agreement explicitly prohibits training on your inputs. Default to enterprise-licensed tools with data processing agreements rather than consumer accounts.

Risk 2: AI-Generated Errors Polluting Your Decisions

AI tools produce outputs that look authoritative — structured, well-formatted, confident. When those outputs are incorrect and flow unchecked into your processes, the error does not stay contained. A demand forecast that is systematically off skews your procurement. A contract summary that misses a key clause gets acted on. A lead score built on a flawed model directs your sales team toward the wrong accounts for months before anyone notices.

The risk is not that AI produces errors — every system does. The risk is that AI produces errors at scale, with a surface appearance of reliability that discourages the scrutiny that would catch them.

Mitigations: Treat AI outputs as a first draft requiring human review, calibrated to the cost of being wrong. High-stakes outputs — financial projections, legal documents, strategic analyses — require verification against primary sources. Build feedback loops that flag when AI outputs are regularly being corrected, as this is an early signal that the underlying model or data needs attention.

Risk 3: Runaway Costs

The most capable AI systems are priced on consumption — you pay per query, per document processed, per response generated. In a pilot, this is negligible. At scale, it compounds quickly and in ways that are not always visible until the bill arrives. A system handling thousands of queries daily, processing large document sets, or running continuous analysis can generate costs that bear no relationship to what the pilot suggested.

The fix is not to avoid AI. It is to treat cost architecture as a design requirement from the start, not an afterthought.

Mitigations: Calculate cost-per-query at your expected volume before committing to a design. Use smaller, cheaper models where the task does not require the most powerful systems — most do not. Implement caching for repeated queries. Set hard spend limits during rollout. Assign clear ownership of the cost dashboard to whoever owns the system.

Risk 4: Privacy and Regulatory Exposure

AI systems built on customer data create privacy exposure that did not exist before — not through malice, but through the nature of how they work. A model trained on customer behavior may make inferences that feel intrusive or cross contextual boundaries customers did not consent to. Data collected for one purpose gets used for another. Training datasets retain personal data longer than your retention policy permits.

India's DPDPA creates enforceable obligations here. For businesses with European customers, GDPR applies with penalties of up to 4% of global annual revenue for serious violations.

Mitigations: Use only the data you actually need for each specific purpose. Apply your data retention policies to training datasets, not just operational systems. Confirm in writing that AI vendors do not train on your customer data. Treat privacy review as part of system design, not a sign-off at the end.

Risk 5: AI Deployed on the Wrong Tasks

AI/ML is not a single technology. As covered earlier in this document, there are meaningfully different approaches — classical ML models, large language models, computer vision, and others — each suited to different types of problems. The mistake that consistently leads to poor outcomes is treating them as interchangeable.

An LLM is the right tool for drafting, summarising, and reasoning through language. It is the wrong tool for predicting next month's demand, detecting fraudulent transactions, or classifying defect rates in product imagery — tasks that require pattern recognition across structured data or visual inputs, not language generation. Deploying an LLM for demand forecasting because it is the most visible AI technology available is like using a spreadsheet to run a customer satisfaction survey — the tool is not wrong in general, it is wrong for this job.

The inverse error is equally common: reaching for a simple rules-based system or a classical ML model when the task genuinely requires the reasoning and language capabilities of an LLM and then concluding that "AI didn't work" when the result is poor.

Mitigations: Before selecting a tool, define the problem precisely — what are the inputs, what is the desired output, and what type of reasoning does the task actually require? Match the approach to the problem, not to what is most familiar or most discussed. This is one of the less visible but more consequential decisions in any AI implementation, and it is where domain expertise in both the technology and the business problem earns its value.

Risk 6: Model Drift

A model trained on historical data learns the patterns that existed up to the point the training ended. It does not know when those patterns stop being valid. A forecasting model built during a period of stable demand will continue producing confident-looking forecasts after a competitor enters your market or a supply disruption changes your cost structure — nothing in the model flags that its assumptions no longer hold. The degradation is gradual, the outputs remain plausible-looking, and the problem only surfaces when someone notices that decisions made on those outputs have been consistently wrong.

Mitigations: Model monitoring is a permanent operational responsibility, not a launch-and-forget activity. Define performance metrics for every AI system and track them continuously. Set thresholds that trigger review when accuracy deteriorates. Plan for regular retraining — do not wait for a problem to make it necessary.

Risk 7: Shadow AI

Your employees are already using AI tools independently — drafting communications, summarizing documents, analyzing data. This is happening regardless of whether your organization has a formal AI policy, and it is not inherently a problem. The risk is that it is happening without data policies, without quality controls, and without any visibility into what is being shared with external systems or acted on from unverified outputs.

Mitigations: Acknowledge shadow AI rather than attempting to prohibit it. Create a practical acceptable-use policy covering which data categories may and may not be used with external tools. Provide approved tools with appropriate data agreements so employees do not default to personal accounts. Build enough AI literacy in your team that people understand when an output needs verification before it is acted on.

Risk 8: Vendor Lock-In

The AI tools market is moving fast. The leading model today is regularly displaced within months. Pricing changes. Providers are acquired. Businesses that build critical processes directly on a single provider's infrastructure — without abstraction between their business logic and the underlying technology — find themselves unable to switch when circumstances change. More insidiously, businesses that outsource their AI strategy entirely to one vendor also outsource their understanding of how their own systems work. When something goes wrong, they lack the institutional knowledge to diagnose it.

Mitigations: Build abstraction layers between your business logic and specific AI providers. Retain ownership of your training data and performance records. Ensure that the understanding of how your AI systems work lives inside your organization, not exclusively with a vendor. Validate performance at real scale before committing to long-term contracts.

The Common Thread

Every risk described above has the same underlying cause: a mismatch between what an AI system was designed to do and the context in which it is actually being used. The pilot worked; production was different. The model was accurate; the world changed. The tool was built for ideation; it was trusted with compliance.
Closing these gaps requires someone who understands both the technology and the business deeply enough to see where they diverge. Getting the design right from the beginning is substantially cheaper than correcting it after the system is live. A data leak cannot be undone. A cost structure embedded in a customer-facing product is hard to change without disruption. Privacy violations discovered in an audit create liability that policy changes cannot retroactively resolve.
The businesses that deploy AI successfully are not the ones that move fastest. They are the ones that design thoughtfully — and that have the right expertise in the room when those design decisions are being made.

Chapter 7

How To Proceed – A Practical Path Forward

This document has covered a lot of ground. The natural question is: given all of this, where do we actually start?\

The honest answer is that the starting point is different for every business, and it cannot be determined without understanding your specific situation: which data systems you have, how clean that data is, which parts of your P&L have the most room to move, where your team's time is currently being consumed by work that AI could handle, and what your organizational capacity to absorb change looks like.

What we know from the pattern of successful AI/ML implementations is this:

- The businesses that start with specific, well-defined problems get better results than those that start with broad AI strategies.
- The businesses that invest in understanding their data before selecting tools save significantly on implementation costs.
- The businesses that design human checkpoints into their AI workflows from the beginning avoid the costly mistakes that happen when AI is trusted with decisions it should not make alone.
- The businesses that measure impact with discipline — comparing before and after with controls — build genuine confidence in what is working and can scale it accordingly.

Why AI/ML, Not Just AI

It is worth returning to one of the opening points. This document uses AI/ML deliberately, because the evidence base for machine learning as a business tool is now more than twenty years old. Google, Meta, and Amazon are trillion-dollar businesses that have leveraged ML as a core competitive capability for most of their existence. The proof is in the P&L, at a scale that leaves no room for doubt.

The more recent wave of generative AI — the LLMs, the image generators, the conversational interfaces — adds new capabilities on top of this foundation. But the application of probabilistic AI to deterministic business decisions is genuinely difficult.

> **The challenge is not building a model. It is ensuring that a system that produces probabilistic outputs maps reliably onto business processes that require accountable, repeatable decisions.**

This translation — from probabilistic AI to deterministic business — is where most implementations stumble. It requires someone who understands both the technology and the business at a level of depth that goes beyond either alone. A pure technologist will build you a model that works in testing and fails in production because it was not designed around real business constraints. A pure business generalist will select a tool that solves the problem they can see and miss the data issues that will undermine it.

What We Would Like to Do Together

If you have read this far and find yourself thinking that there are two or three specific areas where AI/ML could meaningfully impact your business — we would like to understand those problems with you directly.

Our approach is to start by mapping your workflows: how data actually moves through your operations, where decisions are being made, where the highest-volume repetitive work is occurring, and where your current data infrastructure is solid versus where it needs attention. This is a diagnostic process, and the output is a specific, prioritized view of where the highest-leverage AI/ML opportunities are for your business, given your actual constraints.

From there, we design and implement in phases — starting with the highest-confidence, lowest-risk applications that can demonstrate measurable value quickly and using those wins to build the organizational comfort and data infrastructure for more ambitious applications over time.

The goal is not to implement the most impressive AI/ML system. The goal is to implement the one that moves your P&L in ways that are real, measurable, and sustainable. Those are sometimes the same thing. Often they are not.

A Starting Point

We suggest beginning with a workflow mapping conversation — two to three hours with your key operations leads — to identify the three or four areas where AI/ML has the clearest path to impact, given your current data situation. From there, we can scope a focused pilot that produces results you can measure within 60 to 90 days. No long-term commitment required at the outset. The results should speak for themselves.

For questions or to schedule a workflow mapping conversation, please reach out to us at hello@aegion.so.

[1] https://www.nber.org/papers/w31161
[2] https://channelx.world/2023/07/zalando-body-measurement-feature-reduces-returns-10/
[3] https://www.science.org/doi/10.1126/science.adh2586