

© 2026 Alex Kolokolov, Alisa Byshovets  
All rights reserved.

No part of this book may be reproduced, stored, or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without prior written permission of the authors, except for brief quotations used in reviews or scholarly analysis.

*Author:*

*Alex Kolokolov*

*Co-author:*

*Alisa Byshovets*

*Editor:*

*Ksenia Koroleva*

*Illustrations:*

*Alisa Byshovets*

Some illustrations in this book (excluding screenshots) were partially generated using various artificial intelligence tools for demonstration and educational purposes. AI tools mentioned in this book evolve rapidly. Examples and tools are accurate at the time of writing but may change over time. This book is published independently.

HUMAN + AI SYNERGY.  
REDEFINING BUSINESS  
PROCESSES

## *Legal Notice*

*This book is intended for informational and educational purposes only. The content provided herein does not constitute professional, legal, business, or technical advice. Readers are solely responsible for evaluating and applying any information or techniques described in this book to their own use cases, environments, and jurisdictions.*

*All names, companies, case studies, examples, and scenarios included in the book are either fictional or used for illustrative purposes only. Any resemblance to real persons (living or dead), organizations, or actual events is purely coincidental. Exceptions are made for instances explicitly identified by the author as personal experiences: such episodes are based on real events from the author's life. However, names and other identifying details of third parties mentioned in these stories have been changed or omitted to protect their privacy*

*The authors and publishers do not recommend or encourage the uploading, processing, or sharing of confidential, sensitive, or personally identifiable information (PII) in any AI tools, platforms, or services mentioned. It is the sole responsibility of the reader to ensure compliance with all applicable laws and regulations, including but not limited to data protection, cybersecurity, nondisclosure, and export control laws.*

*Under no circumstances shall the authors, editors, contributors, or publishers be held liable for any loss, damage, legal claim, or consequence resulting from the use or misuse of any information, code, prompts, or advice presented in this book.*

*Use your judgment. Seek professional counsel where appropriate. You are in full control—and solely responsible—for any decisions or actions taken based on this content.*

# TABLE OF CONTENTS

Introduction	10
<b>WHO THIS BOOK IS FOR.....</b>	<b>12</b>
<b><u>FROM THE AUTHOR. A LOOK BACK.....</u></b>	<b><u>18</u></b>
THE 2000s.....	20
THE 2010s.....	22
THE 2020s.....	25
<b><u>PART 1. PREPARING LEADERS FOR WORKING WITH AI ....</u></b>	<b><u>29</u></b>
<b><u>CHAPTER 1. THE AI INFERIORITY COMPLEX .....</u></b>	<b><u>30</u></b>
AI-NALYTICS .....	33
BUS-AI-NESS PROCESSES .....	36
WHAT TO DO? .....	38
6 KEY AREAS OF AI COMPETENCIES .....	39
INFORMATION RESEARCH.....	40
DOCUMENT MANAGEMENT .....	41
DATA ANALYTICS .....	42
DATA VISUALIZATION .....	44
AUDIO AND VIDEO PROCESSING.....	45
WORKFLOW AUTOMATION .....	46
<b>THE THINKING CYCLE AND HOW TO SPIN IT UP WITH AI .....</b>	<b>48</b>
THE ASP THINKING CYCLE.....	50
HOW AI SPEEDS UP THE THINKING CYCLE.....	52
REACHING THE HUMAN + AI SYNERGY .....	54

## **CHAPTER 2. TECH TRENDS AND THE MARKET LANDSCAPE**.....55

<b>THE AI ARMS RACE</b> .....	<b>55</b>
THE INTELLIGENCE VECTOR: ANI → AGI → ASI .....	55
THE WEAPONS VECTOR .....	58
THE MARKET VECTOR .....	60
THE ENERGY EFFICIENCY VECTOR.....	63
<b>MEET THE AI EMPLOYEES TEAM</b> .....	<b>66</b>
CHATGPT — JACK OF ALL TRADES (AND GLITCHES) .....	68
CLAUDE — THE SERIOUS ANALYST, SHORT ON EMPATHY .....	70
DEEPSEEK — THE TIRELESS MATH MIND FROM CHINA .....	72
GROK — FAST, BOLD, AND A BIT RECKLESS (BORN IN MUSK’S ORBIT) .....	73
MICROSOFT COPILOT — THE SOLID MIDDLE MANAGER FOR MICROSOFT FANS .....	74
GEMINI — GOOGLE’S BRAINY KNOW-IT-ALL.....	76
DALL-E — THE DYSLEXIC ILLUSTRATOR FOR BUSINESS GRAPHICS ....	80

## **CHAPTER 3. INTEGRATING AI INTO COMPANY CULTURE** .82

<b>THE AI MATURITY MODEL</b> .....	<b>82</b>
THE CHILDHOOD STAGE .....	82
THE COLLEGE STAGE.....	84
THE ADULT STAGE.....	85
THE MASTERY STAGE.....	86
<b>HOW TO START AN AI PROJECT</b> .....	<b>88</b>
RESULT DESCRIPTION.....	89
PROJECT PLAN .....	92
RESOURCES.....	93
RISKS .....	94
<b>ETHICAL ASPECTS OF WORKING WITH AI</b> .....	<b>97</b>
WHO WE’RE RESPONSIBLE TO WHEN USING AI .....	97

**CHAPTER 4. PRACTICE: A VIRTUAL BOARD OF DIRECTORS**  
**..... 104**

**1. DEFINING THE PROBLEM ..... 105**  
**2. CASTING YOUR BOARD OF DIRECTORS ..... 108**  
**3. BRAINSTORMING ROUNDS..... 109**  
**4. SELECTING AND DEVELOPING AN IDEA ..... 110**  
**5. ACTION PLAN AND MEETING PROTOCOL ..... 111**

**PART 2. PRACTICAL TOOLS FOR WORKING WITH AI**  
**SERVICES ..... 113**

**CHAPTER 5. WORKING WITH DATA..... 113**

**DATA CLEANING AND TRANSFORMATION..... 115**  
**COMBINING MULTIPLE FILES..... 120**  
**DATA RESEARCH ..... 124**  
**DATA CONFIDENTIALITY ..... 131**  
**CONCLUSION..... 135**  
**FINAL TIPS FOR WORKING WITH DATA IN AI ..... 135**  
**TOOLS..... 136**  
**CHAPTER 5 SUMMARY ..... 138**

**CHAPTER 6. WORKING WITH DOCUMENTS ..... 140**

**ANALYTICAL DOCUMENTS ..... 142**  
**MAIN STEPS FOR CREATING AN ANALYTICAL DOCUMENT:..... 142**  
**REGULATORY DOCUMENTS..... 151**  
**HOW TO SET UP AI FOR LEGAL TASKS ..... 155**  
**CREATING REGULATIONS WITH AI ..... 158**  
**TECHNICAL SPECIFICATIONS WITH AI..... 161**  
**SALES DOCUMENTS ..... 162**  
**METHODOLOGICAL FRAMEWORK FOR SALES TEXTS..... 166**

WHAT ABOUT PRESENTATIONS? .....	170
AI INSTEAD OF A CREATIVE AGENCY .....	173
<b>CONCLUSION.....</b>	<b>177</b>
<b>CHAPTER 6 SUMMARY .....</b>	<b>179</b>

**CHAPTER 7. VISUALIZATION ..... 181**

<b>DASHBOARDS AND CHARTS .....</b>	<b>183</b>
GENERAL DASHBOARD STRUCTURE .....	186
CONFIGURING CHARTS.....	190
KPI CARDS.....	194
FILTERS AND INTERACTIVITY .....	197
<b>BUSINESS INFOGRAPHICS .....</b>	<b>200</b>
NAPKIN.AI – SCHEMATICS YOU CAN EDIT.....	209
APPROACH TO CREATING INFOGRAPHICS WITH AI .....	212
<b>CONCLUSION.....</b>	<b>216</b>
<b>CHAPTER 7 SUMMARY .....</b>	<b>217</b>

**CHAPTER 8. BUSINESS PROCESS AUTOMATION..... 218**

<b>TERMINOLOGY .....</b>	<b>221</b>
ROBOTS AND BOTS.....	221
HOW A BOT WORKS.....	222
ROBOTIC PROCESS AUTOMATION (RPA).....	223
HOW RPA WORKS.....	224
TYPES OF RPA ROBOTS.....	225
RPA ARCHITECTURE .....	226
EXAMPLES OF PRACTICAL USE .....	227
LIMITATIONS AND VULNERABILITIES OF RPA.....	228
<b>AGENTS AND ASSISTANTS.....</b>	<b>229</b>
ASSISTANTS.....	231
AGENTS.....	233
EXTERNAL AGENT FUNCTIONS.....	236
<b>TOOLS .....</b>	<b>238</b>
SIMPLE BOT BUILDERS .....	238

VISUAL INTEGRATION PLATFORMS.....	246
<b>RPA PLATFORMS .....</b>	<b>256</b>
TOP 3 GLOBAL RPA PLATFORMS.....	258
<b>PRACTICE. AUTOMATION TEMPLATE FOR COMPETITOR MONITORING</b> <b>.....</b>	<b>263</b>
STEP 1. SETTING UP THE MAKE SCENARIO.....	265
STEP 2. ADDING TELEGRAM CHANNELS.....	267
STEP 3. MONITORING JOB POSTINGS.....	268
STEP 4. TESTING.....	269
<b>CHAPTER 8 SUMMARY .....</b>	<b>270</b>
<b>CONCLUSION.....</b>	<b>272</b>

**PART 3. WORKSHOP: CASE STUDIES ANALYSIS.....274**

**CHAPTER 9. BONUS CALCULATOR FOR EMPLOYEES..... 277**

INITIAL DATA .....	277
DATA PREPARATION AND TRANSFORMATION.....	279
CREATING THE CALCULATOR .....	280
FINAL VERSION.....	284

**CHAPTER 10. FINANCIAL DASHBOARD .....286**

INITIAL DATA .....	286
PREPARING AND TRANSFORMING INITIAL DATA.....	288
CREATING MEASURES .....	294
CREATING THE DASHBOARD.....	297

**CHAPTER 11. BUSINESS PLAN .....304**

CLASSICAL BUSINESS PLANNING MODEL .....	306
TRADITIONAL BUSINESS PLANNING.....	306

DIGITAL BUSINESS PLANNING (UNIT ECONOMICS AND DIGITAL PRODUCTS).....	308
<b>CREATING A BUSINESS PLAN .....</b>	<b>312</b>
<b>CREATING A PRESENTATION.....</b>	<b>319</b>
STEP 1. LOADING DATA INTO AI .....	319
STEP 2. ANALYZE AND CLARIFY.....	321
STEP 3. A PINCH OF CREATIVITY (OPTIONAL).....	321
STEP 4. ASSEMBLE THE DOCUMENT STRUCTURE AND UPLOAD TO GAMMA APP.....	322
STEP 5. POLISH THE RESULT .....	325
<b>CONCLUSIONS .....</b>	<b>328</b>

**CONCLUSION. HOW BUSINESSES AND LEADERS CAN AVOID THE AI TRAP..... 330**

<b>TOP 5 BUSINESS DOUBTS AND OBJECTIONS WHEN IMPLEMENTING NEURAL NETWORKS .....</b>	<b>331</b>
<b>TOP 5 OPPORTUNITIES AND TRENDS (DESPITE THE FEARS AND FAILED CASES ABOVE) .....</b>	<b>337</b>
<b>CONSCIOUS USE .....</b>	<b>342</b>
<b>WHAT'S NEXT? .....</b>	<b>344</b>

# Introduction

The story of artificial intelligence today feels a lot like the Tower of Babel. According to the biblical tale, people once set out on an audacious project: to build a tower reaching the heavens, a monument to their own power. But they were punished — languages were scrambled, and suddenly no one could understand each other. Construction stopped and humanity scattered across the earth.

That's pretty much what's happening with AI. Some executives talk only about investments and ROI, while others worry about fines for mishandling personal data or breaking privacy laws. Tech specialists throw around jargon about tokens and billion-parameter models. Meanwhile, enthusiasts debate the coming singularity and the rise of superintelligence—yet in everyday use, the average person just sees a chatbot hallucinate on tasks that require even a pinch of common sense.

The market is overflowing with AI tools and companies promising “AI employees,” cost savings, and revenue growth. But in reality many AI projects only look good

in pilot tests—when it comes to scaling, they often lead to financial losses and reputational damage. And yet, strategically, it's clear: you can't sit this one out. Without investing in AI, you risk falling behind the competition.

With this book I want to build a common language—a way to sync people's expectations. I want to show how managers and specialists can align on goals so that AI stops being a black box. The point isn't to replace humans with machines, but to connect their strengths. Human intelligence stays in charge: it sets direction, makes decisions, and manages risk. Machine intelligence amplifies it—automating the grind, accelerating calculations, and expanding horizons. And when the two work together, you get synergy—the kind that delivers real returns for business.

## Who this Book is For

This book is for people who work with people—and who make decisions. It's for those who've already heard about AI, maybe even tried it, but didn't get a clear or predictable result. It's for anyone who wants to figure out how to plug AI into today's business processes, and how to prepare for the technology and market shifts coming our way.

### **Mid-level managers**

Your bosses expect you to make a breakthrough with AI: boost efficiency, cut costs, speed things up. But you don't have a dedicated budget for experiments, no developer team, no IT support. Deep down, you know: if you don't figure it out yourself, nothing will change.

### **Owners of small and mid-sized businesses**

You're used to counting money and measuring results. Vendors pitch you dozens of solutions, promising “digital employees” and a “full AI transformation.” But you want to understand the principles, see real examples, and know where investment makes sense—and where it's just a waste of time and money.

## Specialists who want to grow

You've dabbled—asked chatbots everyday questions, used them for ideas or quick text drafts. But how to actually use AI at work and get precise, repeatable results? That part's still fuzzy. And it matters, because your career is moving forward—new roles, promotions, tougher tasks. The truth is, people who already know how to work with AI are getting ahead.

Does this sound like you?

- You've tried AI once, but it flopped. Maybe the result was useless, or just not worth the effort. So for now, you're sticking with your old tools.
- Accuracy matters to you. You've seen how chatbots hallucinate and make things up. But when it comes to documents, reports, or technical specs—that's a hard no.
- You value your time. You don't have the bandwidth to chase every update, watch webinars, or run endless experiments. Your day is already packed with sales, production, logistics, people management, and finances.

And yet—you know it's important to figure out how humans and machines can work together. Because only then will you get stable results, without the financial or reputational risks.

### **Who This Book Is *Not* For**

If you're already a technical expert—comfortable with different models, working with APIs, writing bots — this book probably won't blow your mind.

The only thing I'll point out is the environment you work in. If your clients, bosses, or teammates expect AI to do the impossible—or still don't support your projects—hand them this book. It might help you get on the same page.

### **What Makes This Book Unique**

I've picked up plenty of books on AI, but they were often written in over-wordy style that is long on words but short on ideas: filled with survey results and executive interviews, all backing up the same message — “AI is for everyone, and it's going to change everything soon.” Lots of hype and motivation, but very little about what to actually do with the information. No clear path on how to implement it, or even where to start.

Then there are the flashy titles that turn out to be purely technical books on machine learning—aimed at programmers, algorithms, and code, not business problems.

On the American market, I even came across books supposedly written for managers. The problem? You could tell they'd been generated by AI, just recycling American motivational books with some “scientists have confirmed what your mom always told you” thrown in for good measure.

That's why I can proudly say: this book was written by a human.

I'm drawing on 15 years of experience implementing corporate reporting systems. Why does that matter for AI? Because during those years I had to act as a translator—turning business language into technical terms, and dealing with clients who expected miracles from analytics and forecasting. They thought the BI system would crunch their data and simply spit out decisions. Sound familiar? It's exactly the same story with AI.

If you've read any of my earlier books ([Make Your Data Speak](#), [Data Visualization with Power BI](#)), you already know my style. Instead of bland generalizations—which

AI tools are so good at churning out—I share my own opinions and give you concrete, real-world cases from practice.

## **What You'll Get Out of This**

You'll figure out which business problems AI can actually help solve. You'll learn the specifics of working with different types of documents and processes. You'll see how to combine multiple tools into one workflow. You'll understand how an iterative approach to AI works—how to fix mistakes, handle hallucinations, and gradually “train” your own AI assistant.

The book is packed with real cases—stories of both successes and failures. They'll give you perspective, help you spot patterns, and let you learn from other people's mistakes. You'll also walk away with clear terminology so you can speak the same language with colleagues and actually get support for your initiatives.

In the end, instead of blind enthusiasm or knee-jerk skepticism, you'll develop your own informed opinion: how AI can be used in your organization, how it should be controlled, and how it can be built into the way people already work.

## **What You Won't Get**

You won't find a magic prompt here that instantly gives you profit with zero risk and zero cost. And even if you follow the approaches in this book, I can't promise they'll work perfectly in your specific project.

AI is inherently variable—that's just how it works. Learning to deal with that variability is part of the game. But that's also where the real upside is: the synergy you can create once you start applying AI in your work.

So — good luck on this journey. It's going to be a fascinating one.

# FROM THE AUTHOR. A LOOK BACK

Artificial intelligence loves to generalize. The problem is details get lost and those empty spaces in reports or documents get filled with “industry averages” and vague market trends. The essence disappears. That’s why I honestly believe that today, a subjective point of view—an author’s perspective—is more valuable than ever. To form your own opinion on anything, you need multiple perspectives, not just a neat summary from ChatGPT.

I’ve been working with data analytics projects since 2009. First as an employee, then as the head of my own company—personally responsible for results delivered to clients. Most of the key projects were in heavy industries: steelmaking, oil & gas, rail transport. And I know firsthand the gap between how executives talk about “digital transformation” and “innovation” in glass towers in downtown Manhattan, and how things actually look on the ground.

In this opening chapter I want to share how I've watched machine learning and artificial intelligence evolve—through the lens of my own experience. To my surprise, it already spans three decades:

- The 2000s — what I call the era of IT adoption and business process automation
- The 2010s — the era of big data and the golden age of analytics
- The 2020s — the era of artificial intelligence everywhere

I'm not claiming objectivity here. This is just one participant's perspective.

## The 2000s

I remember 2005: in my junior year at the University, we had a course on “Intelligent Information Systems,” taught by Professor Goldstein. What struck me was the idea that you didn’t have to hard-code algorithms—you could build a system that mapped relationships on its own, like neurons in a brain, and found the best way to solve a problem.

The phrase “**cluster analysis**” sounded especially beautiful at the time. The concept itself is simple: each object has a set of attributes—say, for a customer, it might be purchase history, visit frequency, and average spend. The algorithm measures the “distance” between objects across those attributes and groups similar ones together. Instead of a plain segmentation by income bands or demographics, you suddenly get behavioural clusters that aren’t obvious at first glance.

**Decision trees** stuck with me too. The algorithm looks for yes/no questions to split the data. For example: “Did the customer buy something last month?” “Is the average purchase over fifty dollars?” At each step, the tree branches out, creating more specific conditions. Follow the branches far enough, and you end up with predictions: Will the customer churn? Will they buy again? Should we approve this application?

Later, in grad school, I dug deeper into the subject—and realized it was rarely applied in practice. Sure, there were projects predicting state or city-level economic growth, but they didn't have much to do with reality. To get reliable results, you need huge datasets. And back then, most companies only had basic automation—some accounting software, maybe a primitive CRM.

The first real adopters were banks and large retail chains, where transactions happen every second. But even there, forecasts ignored outside factors, so artificial intelligence projects were often left in the hands of math romantics dreaming of algorithmic decision-making—or pragmatic managers spending innovation budgets, fully aware their projects would just sit in a drawer.

## The 2010s

The next decade changed everything. Two factors gave predictive technologies a real push:

- First, companies had automated most of their core processes and built up years of data—three, five, sometimes more. That meant forecasts could now account for seasonality and other trends.
- Second, server power got dramatically cheaper. Adding a hundred gigs of RAM or a terabyte of storage was suddenly affordable not just for big corporations but also for mid-sized companies. On top of that, cloud data centers made it easy to rent resources on demand.

At that time, my company's main focus was building corporate data warehouses and automating reporting. And suddenly my clients were getting more than just pretty charts on dashboards—they could drill all the way down into detailed data on customers, products, and employees. That naturally led to the big question: “What if...?” They wanted to model different scenarios, to predict how customers would behave.

Our first real project came in 2014: predicting churn for a regional internet provider. The idea was to build a “profile of a leaving customer,” spot those most likely to quit ahead of time, and then run retention campaigns.

Back then we used IBM SPSS. We loaded in customer data: service usage activity, support tickets, and other factors. We left out the most recent quarter—so we could later test how accurate the prediction was. At first, the error rate was pretty high. But after another month, we retrained the model with fresh data and added competitor stats (the client finally had a reason to clean that up).

The prediction accuracy hit 80%, which was considered solid back then. We proudly went to present it to the commercial director. His reaction? “Thanks, guys. But I already know our main churn segment—it’s the suburban communities. Our competitor is offering them more services at half the price. I was hoping you’d predict where we can find new customers...”

That was my wake-up call. For mathematicians success means prediction accuracy (and not just one metric, but a whole set of them). For business? All those smart words were basically useless. Most clients—except for a rare few—wanted fast results without putting in much effort themselves.

So I shifted focus to dashboards, where clients were actually ready and demand kept growing.

Over the years, companies collected more and more data, and a new trend started to solidify: Data Science. Underneath, it was still the same mathematical statistics, just with more modern tools. And the job market quickly filled up with specialists who could use them.

By 2017–2018, I decided once again to test in practice whether a machine could really think for a human.

A plant producing industrial resins. The task: predict the resin's output parameters from the reactor. Why? Because those parameters determined where the batch could go—either to the wood-processing industry (say, for making cabinets) or to tire manufacturing. The problem was that orders might come in “for wood,” but the output sometimes turned out fit only “for rubber.”

The client swore they had all the data, with a fully automated lab. In reality, everything hit a wall at data collection. Turns out, some measurements were still being logged in paper journals. And to make it worse, a technician could leave early and just jot down averaged numbers. Which meant real deviations from the production process weren't being tracked at all.

We had to build a custom app for entering data. But new barriers appeared right away. The data scientists kept talking in terms of variation coefficients without really understanding the production itself. To move things

forward, we had to bring in a PhD chemist—someone who could challenge the client’s process engineers on an expert level. Only that kind of authority could break through the resistance of the old-school conservatives. In the end, the project wasn’t really about machine learning at all. It turned into complex consulting—basically, “moving in with the client.” And those so-called “data scientists”? They were still just programmers and mathematicians. We were still far from a replicable, conveyor-belt model of projects.

## The 2020s

As the years went by clients had more and more data—and more areas to apply it. One new frontier was text analysis and extracting meaning.

In 2021 we worked with an industrial holding company with over 100,000 employees. We were doing HR analytics, including employee satisfaction studies. Beyond the eNPS score, there were thousands of open-text comments—the real gold, because they carried the nuance you couldn’t see in just numbers.

But a million comments can’t be processed manually. So we built a model that would categorize comments and detect sentiment—positive or negative.

First, the client manually classified the first thousand responses, then we trained the model. Accuracy came out at 95%, which hit the technical target. But when we

put the results on a dashboard and showed them to a focus group, it was almost a flop. The clustering errors always seemed to land right on the first screens or after just one or two filters. A complaint about salary got thrown into “working conditions.” A note about overtime ended up under “internal policy.” To the users, it felt like the accuracy was closer to 50%.

We knew exactly what to fix, how to retrain the model. But we had two days, not two weeks. So we exported 10,000 problematic records and brute-forced it: two days of manual cleanup in emergency mode. After the presentation, we retrained the model properly so new data would be interpreted correctly. But it left a bad taste in my mouth

Fast-forward to today: we’ve got practically unlimited computing power sitting in our pockets. What used to take weeks or months to train now takes hours—or even minutes. From idea to working prototype can be just a few days. The real bottleneck isn’t the machine anymore—it’s the people. To interpret the results, discuss them, and agree on joint actions still takes days, weeks, or in bureaucratic organizations, months. But the fundamental shift has already happened.

Here's a case. In 2021 we tried to launch a demand forecasting project for a meat producer. Off-the-shelf solutions didn't fit at all. Classical regression models required deep customization to deal with their seasonality quirks and messy accounting specifics. The full budget—including software licenses, custom dev work, and integration—topped \$1M USD. Even though we calculated an ROI many times higher than the costs, the client still decided the risk was too high and put it on hold. They kept calculating reorder points in plain old Excel.

By 2024 the time had finally come. Processes and data had matured enough, so we piloted using the Anthropic API. We went through a few improvement cycles and then ran a strategy session. We brought in leaders and specialists from all departments involved: production, logistics, sales, marketing. During the session, we tested hypotheses live—evaluating scenarios both with data and with expert judgment. The result? Decisions on process changes and budget approvals got made in *one day* instead of three months.

After the pilot, the model was deployed locally on the client's servers and scaled up. The costs were dramatically lower compared to the “previous generation” architecture. And the payoff was clear: the company was able to boost production and sales during peak season by 20% compared to the year before.

Looking back over the last 20 years, it's obvious: the tech keeps changing, the names of approaches keep changing — but the logic stays the same. There have always been overblown promises from IT vendors. And naïve expectations from business leaders that “the IT guys will magically grow profits for us.” But there have also always been those who knew how to use technology wisely in business.

Let's hope this time, you and I make it work.

# PART 1. Preparing leaders for working with AI

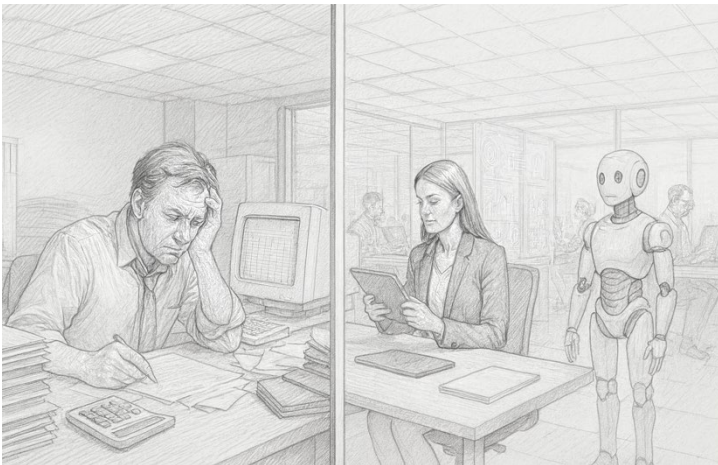
Bringing AI into business processes isn't just about learning new tools—it's about changing the way you work. This part of the book will help you prepare yourself and your team to use artificial intelligence in a way that adds value instead of creating more chaos.

The truth is leaders often swing to extremes. Some overestimate AI, expecting it to solve every problem. Others underestimate it, dismissing it as just a toy. Both approaches lead to dead ends: the first to disappointment, the second to missed opportunities. If you want AI to be your ally, you need to see both its strengths and its limitations — and, above all, keep control in your own hands.

Here, we'll look at things from a strategic angle: which business processes can realistically benefit from artificial intelligence, and what the global trends in AI adoption look like. On the flip side we'll talk about how to develop your team's skills, and how to gauge the maturity level of people and groups when it comes to AI. That will give you clarity on where to focus your efforts—and how to avoid the common traps and misconceptions that derail so many projects.

# CHAPTER 1. The AI Inferiority Complex

Everybody's talking about Artificial Intelligence. First, it was about the idea itself but now for years we've been hearing about how companies have already "hired AI employees," launched marketing and sales with zero budget, and are living the future. It can start to feel like your company is the only one left behind—still doing everything "by hand." And online the chatter never stops about endless lists of prompts that supposedly earn a thousand dollars an hour, save months of work, or let a single person launch an entire business overnight.



That's how the inferiority complex kicks in. You feel embarrassed to admit you're not on trend. So when someone asks "Are you using AI?" you answer, "Of course, we're rolling it out everywhere—it's our top strategic priority!"

Watching the market, I've noticed a quiet shift: the line between "we've implemented" and "we've tried" has blurred. The honest answer most companies could give is: "We tried, but..." Almost everyone has had their first attempt at solving business problems with artificial intelligence—and for 95% of them it failed.

That first failure? Totally normal. We're in a period that's comparable to the industrial revolution, and it's going to take more than one attempt. This shift in both technology and society will play out over the next decade. It's unrealistic to expect a revolutionary breakthrough in one or two years. The companies now bragging about success stories? They started investing in AI five or even ten years ago—long before it was a buzzword.

What isn't normal is living in illusions, holding inflated expectations. And that's the real issue in the market. Service providers keep promising quick wins—that you can just "plug in AI" without involving top management or business owners. That somehow, without even

engaging the IT team, artificial intelligence will write its own code and set everything up turnkey.

Honestly, I don't feel much sympathy for companies that fall for that pitch. My sympathy is for the employees who then have to clean up the mess. I get calls all the time asking for fast results—but usually after someone else failed to deliver. In my experience, there are two common scenarios where expectations for AI fall flat: in analytics, and in business processes.

# AI-nalytics

The idealized vision goes something like this: a “neuro-analyst” who pulls together all your data, organizes it neatly, explains why targets weren’t met (and who’s to blame), and tells you exactly how to hit next month’s goals.

*“We want to build a chatbot on top of our BI system—for executives! So they can ask anything about revenue, profit, expenses and get instant answers, without touching databases or dashboards.”*

*“But weren’t the dashboards in the BI system built for that exact purpose? So executives could see clear visuals and get answers without digging into databases.”*

*“Well... yeah. But they don’t use the dashboards.”*

*“Why not? Missing KPIs? Or what’s the problem?”*

*“We never asked... We just built it based on the agreed scope.”*

To me, this looks exactly like what we saw 10 years ago, 5 years ago—and what many still see today—with BI analytics. Executives don’t use dashboards because:

1. **The right data isn’t there.** A typical scenario: accounting data gets exported from the ERP

system, but management reporting is still stuck in Excel. A dashboard with slick quarterly EBITDA trend charts is of little use if you need to figure out how to reduce this month's cost of goods sold.

2. **They don't want the responsibility.** When all the data is available many middle managers start resisting a transparent analytics system. Why? Because they can no longer say, "I didn't have the data." They're forced to make uncomfortable decisions.

In the end, AI analytics projects run into the exact same wall BI projects did: the data warehouse. Either it doesn't exist at all, or it only has narrow slices of data—finance, sales, inventory. From that you'll never get insights about raw material quality, workforce performance, or the other factors that truly drive competitiveness.

On the surface the hype around "analytics chatbots" looks like it could breathe new life into BI projects—giving leaders an excuse to finally tackle the foundational work that budgets or patience didn't allow for earlier. But here we're back to inflated expectations—only this time, directed at people.

To build a true “centralized” analytics system—from strategic KPIs down to raw transactional data—you need to tie together two key elements:

1. **The KPI tree** (a pyramid of metrics or another hierarchy of business indicators). That requires real immersion in business processes.
2. **The data model.** This means mapping the semantic layer to the physical tables in the warehouse, which takes technical expertise.

And one person can’t realistically do both (except for a rare breed of superstar architects). But of course, companies often dump the whole responsibility on a single specialist. Which leads to two extremes:

- Expecting the analyst, for their salary, to deliver consulting at the level of McKinsey.
- Or IT analysts building a database and telling business users: “It’s self-service now—here’s your Lego set, build whatever you want!”

In this situation, I don’t have any AI hacks—just like I didn’t have any BI hacks. For our clients, building a corporate data warehouse was always a 1–2 year project. And it wasn’t just infrastructure—it also required developing true data-driven skills among managers.

# Bus-AI-ness Processes

It's the same story with automating business processes. Clients expect bots that can talk to customers, understand the full product line, instantly pick the best offers, and handle objections. But it always hits the same wall: there's no solid knowledge base to train artificial intelligence on. Best practices haven't been digitized, and the employees who have the real expertise resist—just like they did 20 years ago when companies were first automating processes and rolling out ERP and CRM systems.

As with analytics, successful AI implementation often means taking a step back—returning to process automation first. The biggest potential is around CRM. Even small and mid-sized businesses have CRM systems, but they mostly use them for the sales funnel: tracking deals, monitoring quotas, automating paperwork. Within that paradigm, everyone obsesses over attribution: where did the client come from, how do we track the lead source reliably, and which ad channel deserves the budget. But the “funnel view” covers only a fraction of what matters about customer behaviour. Once the deal is closed, systematic data collection

stops—no tracking of customer experience, repeat purchase factors, or referrals.

The second most common area is HR. Automated candidate screening, virtual recruiters for first-round interviews—these have been around for a while, with plenty of IT solutions on the market. They're not cheap and usually pay off only for high-volume roles. For specialized hires, recommendations still work better.

But the real pain point isn't recruiting—it's retention. This is where AI could shine: building personalized development plans, matching training formats to individual learning styles, spotting early signs of burnout, and developing the leadership pipeline. But again, the problem is lack of structured data. Payroll, training records, and basic demographics don't give enough insight. And mapping out the actual professional journey still belongs to no one. HR doesn't take it on (they lack the technical chops), and IT won't touch it (they're always swamped with "more important" priorities).

## What to Do?

I understand the urge to keep up with the times—and my advice to take a step back probably doesn't sound very inspiring. It feels like if we focus on cleaning up mundane operational issues right now, our competitors will blast off into space on an AI-powered rocket while we're left behind.

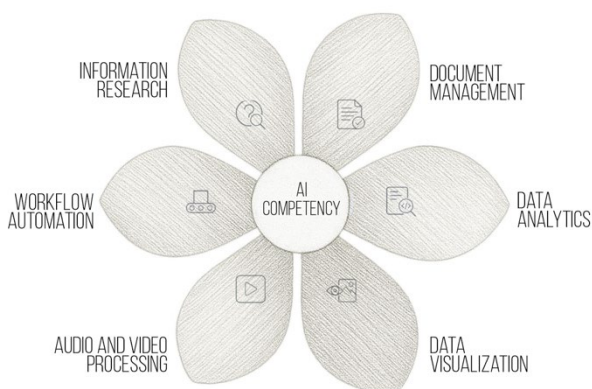
But trust me—most of your competitors are facing the exact same problem: they lack the solid foundation needed to actually implement AI effectively. Even if they do “launch,” they're unlikely to make it past the stratosphere.

On the other hand, if you streamline your business processes and get your analytics in order, you'll be ready to ride any AI wave that comes along—especially since new tools and technologies emerge every month.

Just remember: no matter what service providers promise you, no AI can do the hard work of fixing your core operations for you. That part is still entirely up to you.

## 6 key areas of AI competencies

How do you know if you're actually good with artificial intelligence? These days, AI experience is listed as a required skill in digital job postings. But at the same level of competence, some people modestly call themselves beginners, while others boldly claim "AI expert" and sell consulting sessions. From the employer's side, the question is: how do you objectively assess these skills—and how do you know the candidate didn't just generate their "case studies" with ChatGPT?



Integrating AI into everyday work isn't one skill—it's an entire ecosystem of habits. In this chapter we'll go through six key areas of working with artificial intelligence:

1. Information research
2. Document management
3. Data analytics
4. Data visualization
5. Audio and video processing
6. Workflow Automation

## Information Research

ChatGPT has become the household name for large language models and it's still the global leader. But it's already clear it won't be the permanent industry standard like Excel or Salesforce. That's why it's important to know how to find answers across multiple AI tools—and to spot hallucinations.

A useful habit: run the same query through three different models. For example, ChatGPT, DeepSeek, and Claude. Most likely, ChatGPT will give the most helpful, detailed answer. But not always. Sometimes its “magnetic storm” days hit, and DeepSeek comes out more accurate.

What to watch for when picking your go-to AI assistant for daily tasks:

- **Accuracy.** How often does the model “make stuff up”?
- **Substance.** Claude or Qwen might give you a crisp, straight-to-the-point answer, but it can

feel too dry. ChatGPT might deliver a whole essay with examples—sometimes useful, sometimes overkill.

- **Speed.** GPT-5 can take its sweet time “thinking” for a minute or two. Maybe you need quick back-and-forth, not deep reflection.
- **Context memory.** How well does it remember details in long conversations? Can it handle typos or clumsy prompts and still keep track of the discussion?
- **Censorship.** ChatGPT and Claude are the strictest about ethics—sometimes excessively so. Grok and Mistral are more relaxed. And Perplexity has essentially re-skinned DeepSeek without its political filters.

## Document Management

At the basic level, most people use AI like a search engine. Instead of “just Google it,” we now say “ask ChatGPT.” The next level is document creation.

You’ve probably noticed: AI-generated text has a tell. It often reads like it was written by a mid-level marketer for an average social media post. Business documents

demand more precision, and crafting those with AI takes effort. Like any skill, practice matters. If you've drafted 100 contracts with AI, that 101st one will really take you only a couple of prompts.

The same goes for marketing decks or sales letters: practice helps avoid clichés.

I've been in situations where someone sent me 20+ pages "for approval." On the surface, the text looked convincing: clear structure, examples, sales triggers. But read closely, and it was fluff—just the same ideas rephrased over and over.

"Make me a report in five minutes" is a college student's hack, not a professional's. With AI, you can gather references, pull in case studies, research findings, legal citations, even spin up multiple drafts in different tones. Then you pick the best one. That's how you end up with a genuinely professional, useful document (for details See Chapter 6).

## Data analytics

People love to say artificial intelligence can replace analysts and programmers. But most people's first attempt at analytics with AI ends in disappointment. Yes,

AI makes mistakes—but often the issue is you don't know how to frame the task or guide it away from those mistakes.

Here are typical analytics tasks a modern manager should be able to handle solo, without bothering the team:

- **Merging multiple tables** or files into one dataset, cleaning duplicates.
- **Cleaning values:** removing spaces, standardizing dates and currencies, filling in missing fields.
- **Transformations and groupings:** breaking data down by product, branch, channel, etc.
- **Filtering and anomaly detection:** surfacing your top products, clients, employees, or outliers.
- **Calculated metrics:** building KPIs like ROI, ROAS, LTV, CAC, and double-checking their correctness.
- **Templates for missing data** (vendors, competitors, benchmarks, industry indices).

- **Forecasting and trends:** with enough data, AI can give you at least short-term predictions—a month, a quarter.

In other words, if you're unhappy with AI's analytics, it's probably because you don't yet know how to "talk" to it. With experience you'll find the right approach—and suddenly it feels like you've got a smart assistant working alongside you (for details See Chapter 5).

## Data visualization

At this point, we can safely say AI has learned how to visualize data and create business graphics. The days of awkward "six-fingered" charts are gone. Most AI tools now choose chart types correctly based on the data—but the final output still varies a lot.

- **Text.** Some AI models don't really do visualization. For example, Qwen might promise you a chart, but deliver an ASCII sketch. More placeholder than usable.
- **Images.** Many tools generate a chart as a static image—fine for a presentation, but you can't edit it. True export to Excel or PowerPoint still isn't here.
- **HTML or code files.** This is a solid workaround: AI generates an interactive chart. Update the

data and it refreshes. Drop it on a site, or send the file to a colleague. Claude is especially strong here.

The same principle works for entire dashboards. These can be images or HTML templates where you load data. For small datasets, this can even replace a BI system. Plus, there are specialized AI tools built just for dashboards and visuals: Polymer, Bricks, and others (for details See Chapter 7).

If the first three skills (search, documents, analytics) can be done inside one tool like ChatGPT, professional visualization still requires dedicated tools: diagram builders, infographic makers, presentation apps.

## Audio and Video Processing

The next level is extracting insights from audio and video: call recordings, video conferences.

Computer vision isn't new—it's long been used in specialized software for tasks like quality control on assembly lines, theft detection in retail, or proctoring during online tests. But everyday AI tools aren't quite there yet. Audio, though, they handle well.

The real value is what you do with the transcripts. One powerful area: quality control in sales and call centers. No more listening to hundreds of hours of recordings to spot weak points. AI does the heavy lifting—your job is to improve scripts and train staff.

For meetings with three or more people, you can analyze participation: who contributes ideas, who just sits in. This isn't as widespread as sales QA, and it requires finer tuning. But strategically, it can show who's really doing managerial thinking and who's just "in the room."

As for AI-generated audio or video—that's still more of a novelty. Fun, yes. But leaders have more serious problems to solve.

## Workflow Automation

If you're confident with information, documents, data, visuals, and audio—you already have a serious edge. The next level is automation: building bots for local business tasks.

For this no-code platforms like Make and n8n (pronounced "nate-eight-n"). These tools allow you to create workflows that connect multiple AI services, messaging apps, databases, and documents.

However, using them can come with significant limitations—particularly around data privacy and the ability to scale to industrial-level data volumes. That said, it's essential for any manager to understand the basic principles of bot development. This knowledge will help you set realistic expectations—not only regarding development costs but also in terms of the actual impact these solutions can deliver.

I know entrepreneurs who genuinely run businesses solo now—or rather, with a team of “AI employees”: a CFO, a marketer, a lawyer. They chat with them like real colleagues, even “joke by the water cooler.” But here’s the truth: that wasn’t built in a week. It took a year or more. Still, that’s not a long time for a complete business model transformation.

# The thinking cycle and how to spin it up with AI

A leader's job is making decisions. Analysts can draft a flawless business plan. IT can automate the most complex process. But sooner or later (usually sooner) things go off script: conflicts between people, market shifts, geopolitics. In those moments, there's no time for research or audits. You have to act here and now, with incomplete data—estimating consequences, comparing options, leaning on experience.

Decision-making is one of the brain's most energy-draining functions. That's why it looks for shortcuts—habitual answers—even when those shortcuts lead us into a dead end. And energy is finite: in a day, you can only make so many real decisions.

Think about a “typical” workday—whether in the office or remote. Your door (or your chat app) keeps pinging with questions: approve a contract, confirm a client discount, pick a banner color, check a job post. Nine times out of ten, people could handle it themselves, but the question still lands on your desk. Each quick approval chips away at your mental battery. By lunchtime, it's flashing yellow, and once again you don't have the energy to focus on

strategy, sales forecasts, new products, or critical calls about opening—or closing—branches.

And yet, there are moments when decisions come easy. Maybe in a conversation with a partner, mentor, or consultant. You hear a case study at a conference, talk it through over lunch, figure out how it fits your company, make a few calls, “sell the idea” to loyal clients, and line up an order. The next day, you’ve already drafted a business plan for a new direction.

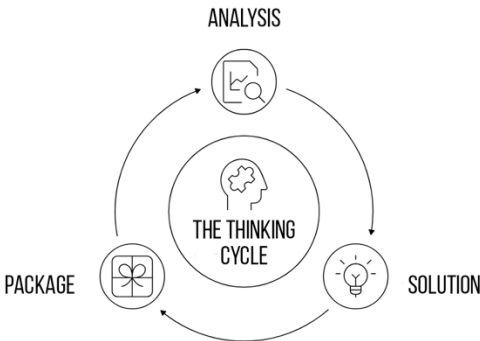
That’s what it looks like when you kick off a full thinking cycle—the kind that sparks breakthrough results in business.

# The ASP Thinking Cycle

That synergy kicks in when someone helps develop your idea in the moment—or when you, in turn, look at someone else’s creative challenge from a systematic angle and suggest solutions. For a solution to be strong, you need to see it from multiple sides and go through three stages:

- Analyze the information
- Search for ideas and options
- Package the result and present it

I call this the ASP thinking cycle (Analysis – Solution – Package).



The truth is most of us lean heavily toward one type of thinking and switching to another mode doesn’t come naturally. Some people excel at analysis and research;

others at generating ideas. And often, we get stuck in our comfort zone without realizing it. You've probably seen analysts who keep saying, "We need to collect more data," instead of making a decision. Or creatives who freeze the moment they need to calculate unit economics.

So to land on a strong solution or drive a strategic project, you either have to push yourself through all three phases—or build a team that covers them. And often, it takes several loops through the cycle.

In my work, analysts often complain about managers: "You ask for a report, then keep changing your mind—add this data, recalculate that scenario. Why didn't you just say it from the start?" The answer is simple: it's impossible to foresee everything in advance. Until you see an interim result, you don't know what's missing. Managerial thinking is iterative by nature.

Every leader dreams of finding a "3-in-1" professional—or at least a "2-in-1." But that's rare.

Take a food company developing a distributor bonus program. They've got the idea, but now they wait for analysts to recalc sales forecasts. Then wait again for marketers to polish the deck, add competitor data, and

sharpen the positioning. Meanwhile, the leader's brain is idling: the task is hanging, but there's still no decision.

The old paradigm of “one person = one role” is no longer competitive. And if you want to evaluate an employee—or the output of a team or contractor—you need to have the competencies yourself. That brings us to the next question: how do you build those competencies with artificial intelligence?

### How AI Speeds Up the Thinking Cycle

Picture this: a “magic button” you can press to hand a task to an analyst, designer, marketer—even a strategy consultant—and get the answer immediately. No waiting days or weeks for “mere mortals” to get the job done.

You've already heard about “AI employees,” the idea that you can delegate routine tasks to AI services. But that magic button jams all the time: the machine thinks like a freshman, and what worked smoothly yesterday may start hallucinating and spitting out nonsense today.

And that's exactly where the potential lies! Magic buttons only exist in fairy tales. While your competitors are out searching for wizards, they're more likely to meet storytellers. You, on the other hand, can invest in developing the skills to work not with one button, but with an entire control panel—and get stable, predictable results.

Here's how it works for me. I'm an analyst but with management experience and exposure to countless client cases, my real value is in generating new ideas, evaluating them, and making decisions. Sure, I can process data and calculate scenarios myself—but only once per session. After that my energy's tapped. Coming up with more ideas or packaging them into project budgets or business plans? Not happening without a break.

Now, working with artificial intelligence, I can keep the cycle going. I throw tasks on the fly to ChatGPT or Claude to gather extra data, then drop the results into Gamma to draft a presentation, or Lovable to mock up a website. With Infogram, I can instantly see what a product catalogue might look like.

Seeing even a rough draft sparks a new round of thinking. Because I'm still in the context, I can spin through two or three full thinking cycles in one sitting. Then I hand my team or contractors a clear assignment. I know exactly how I'll evaluate their work, and I pay only for results.

Our clients find the same thing. For some, the starting point is selling an idea.

Take Michael, commercial director of a logistics company. He can turn “zero into one”—craft a pitch deck, light up clients or partners with sheer charisma. Then he delegates the grunt work to

artificial intelligence: supplier analysis, for example. The next phase is picking options or generating fresh ideas from the data. That gives him fuel to sell the updated idea. On the next turn of the cycle, he hands AI the prep work for tender documentation or the analysis of financial and legal risks.

## Reaching the Human + AI Synergy

Like any mental skill, it takes practice. What I wanted to get across in this chapter is simple: there's another way to do things. Even if you've had a bad first experience with artificial intelligence—even if your expectations weren't met—it's worth giving it another try.

AI technologies will keep coming and going at an even faster pace. But if you rewire your own mental habits to work with them—accepting both their strengths and their flaws—you'll gain a real edge. You'll outperform those who rely only on human intelligence out of habit, and those who naively hand everything over to machines.

# CHAPTER 2. Tech Trends and The Market Landscape

## The AI Arms Race

Leadership in AI isn't just about tech competition—it's about a struggle for global dominance, with stakes higher than the space race or even the Cold War of the 20th century. Countries and corporations are pouring in trillions to gain an edge in the areas shaping the future: from military power to economic control.

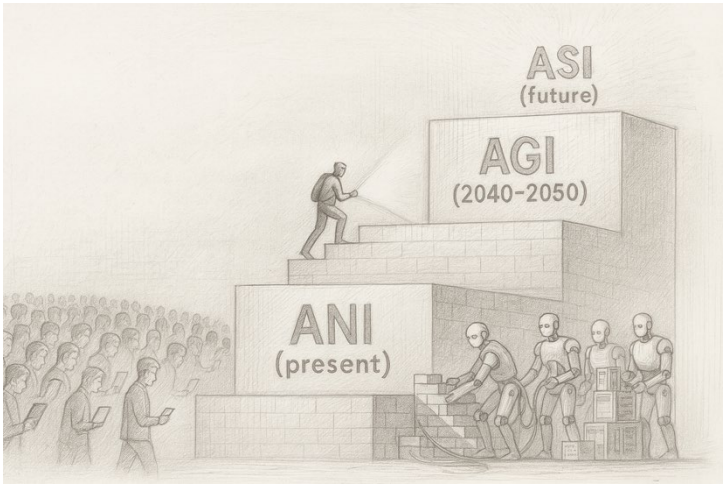
But the goals of this race aren't measured only in money or patents. They're measured in vectors that are reshaping the world order itself. Let's break down the key dimensions of this fight, and why it really does resemble an arms race—where a single breakthrough can flip the balance of power:

1. Intelligence
2. Weaponry
3. Market
4. Energy efficiency

The Intelligence Vector: ANI → AGI → ASI

The ultimate strategic target is creating artificial intelligence that matches—or surpasses—human reasoning. In both academic research (and sci-fi) this path is often described in three stages:

- **ANI (Artificial Narrow Intelligence).** That's what we have today: specialized models like ChatGPT. They can handle specific tasks, sometimes showing breakthroughs (like in coding), but stumble on everyday problems that need basic common sense.
- **AGI (Artificial General Intelligence).** A universal intelligence, able to reason, learn, and tackle a broad range of tasks like a human. AGI implies not specialization, but versatility—adapting to new situations without retraining.
- **ASI (Artificial Super Intelligence).** A superintelligence that outperforms humans across the board and can self-improve, creating newer, better AI faster than we can. ASI is the “point of no return”: it could solve global problems like climate crises—or decide humanity itself is the problem.



The arrival of AGI is forecast somewhere between 2040 and 2050, though optimists like Sam Altman from OpenAI claim it could show up as early as 2030. But nobody wants to come in second in this race.

In January 2025, Donald Trump appeared alongside Larry Ellison (Oracle), Masayoshi Son (SoftBank), and Sam Altman (OpenAI) to announce Project Stargate: “up to \$500 billion” for a network of data centers and “100,000+ new jobs” in the US. It’s both a political gesture and a tech signal: the US is betting big on infrastructure—hardware and energy—to accelerate the path toward AGI<sup>1</sup>.

<sup>1</sup>

<https://www.reuters.com/business/media-telecom/openai-oracle-softbank-plan-five-new-ai-data-centers-500-billion-stargate-2025-09-23/>

## The Weapons Vector

AI is quickly becoming a tool of military power: from autonomous drones and cyberattacks to biotech. The US and China are leading the race, but every country wants its own sovereign stack—its own models, chips, data centers, and supply chains.

China is developing a doctrine of “cognitive warfare”: deploying generative AI systems to process intelligence, spread disinformation against adversaries, support decision-making, and enable early warning. It also showcases integrated “missile–drone–cyber” capabilities in military parades and exercises, keeping the U.S. on edge. In response, the Americans are pouring billions into drone swarms—intelligent, self-organizing UAV networks—and battlefield medicine for rapid medical care in combat zones. Overall, escalation risks are mounting: a single breakthrough in AI could shift the strategic balance as dramatically as nuclear weapons once did.

On the security front, AI is no longer just a tool for defenders — it’s being tested as the attacker itself. In recent research, so-called “agentic AI” systems were set loose in controlled experiments, tasked with scanning websites, chaining together exploits, and writing phishing

lures. The results were sobering: machines that don't sleep, don't get distracted, and don't second-guess themselves can hammer at digital defenses with a relentlessness no human hacker could match. Unlike traditional malware, these agents adapt in real time — trying a hundred variations of an attack in seconds, rewriting their own code on the fly until something cracks open. Security experts warn this is not science fiction but a near-term reality: AI systems behaving as autonomous cyber-offensive units, capable of taking down networks with little more than a prompt<sup>2</sup>.

On March 15, 2025, hackers using AI-generated code broke into California's power grid, causing a 12-hour blackout. The attack exploited vulnerabilities in SCADA systems, while AI-crafted phishing emails—98% indistinguishable from legit messages—tricked their way in. The result: 15 major substations went down, 1.2 million households were hit, and damages topped \$300 million.

The US blamed China and sanctioned three Chinese tech companies suspected of involvement. Insiders say it might have been a test run—probing weaknesses ahead of a possible move on Taiwan.

---

<sup>2</sup> <https://www.anthropic.com/news/detecting-counter-misuse-aug-2025>

# The Market Vector

By 2025, the global AI market is valued at around \$300 billion. Roughly 1.8 billion people use AI in some form, but monetization is lagging: only \$12 billion comes from the consumer side<sup>3</sup>. The real money is in enterprise software—about \$174 billion. Market players aren't just fighting over revenue—they're competing for users. Control of user data means leverage in the global economy.

**OpenAI** leads in revenue (\$12 billion), with the company itself valued at about \$500 billion<sup>4</sup>. Their flagship product, **ChatGPT**, counts around 20 million paying subscribers and 5 million business users through Microsoft Azure integration. That combination of consumer mass adoption and corporate distribution puts OpenAI ahead in both markets.

---

<sup>3</sup> <https://www.abiresearch.com/news-resources/chart-data/report-artificial-intelligence-market-size-global>

<sup>4</sup> <https://www.reuters.com/business/openai-hits-12-billion-annualized-revenue-information-reports-2025-07-31>

**Anthropic**<sup>5</sup>, another US-based company, is booming in the enterprise sector. From 2024 to 2025, its revenue jumped from ~\$1 billion to \$5 billion. The company is valued at \$183 billion, with about 300,000 corporate clients. Each contract brings in far more than a simple OpenAI subscription. Their **Claude** model is embedded in Slack, Notion, and AWS—giving Anthropic steady corporate budgets. If OpenAI makes tens of dollars per user per month, Anthropic makes hundreds of thousands per client annually.

**China's AI players** don't disclose revenue numbers, but their strength is scale. In early 2025, **DeepSeek-V3** claimed its model training cost only \$5.5 million—compared to \$100+ million for American models<sup>6</sup>. Lower costs not only mean cheaper products but also faster rollout across services. Beyond DeepSeek, giants like Alibaba (**Qwen**), Baidu (**Ernie Bot**), and Tencent (**WeChat** ecosystem) are pushing global expansion. Their AI revenue is unclear, but hundreds of millions of users interact daily with their bots, voice assistants, and recommender systems.

---

<sup>5</sup> <https://www.anthropic.com/news/anthropic-raises-series-f-at-usd183b-post-money-valuation>

<sup>6</sup>

**Mistral AI**, Europe’s response to the global race, is smaller: less than \$1 billion in revenue and a valuation of about \$5 billion<sup>7</sup>. But their bet is different: sovereign stacks, open models, and local data centers. For Europe, it’s not just business—it’s politics: owning standards, controlling data, and staying independent from the US and China.

And now—the **top 10 AI startups** worth watching:

Rank	Company	Year Founded	Total Funding (Billion USD)	Company Valuation (Billion USD)	Estimated 2025 Revenue
1	OpenAI	2015	21.9	300	10
2	Anthropic	2021	14.7	183	5
3	xAI	2023	12.1	75	1
4	Databricks	2016	15.25	62	3
5	Mistral AI	2023	2.1	14	0.8
6	Scale AI	2016	1.5	10	0.5
7	Hugging Face	2016	0.4	5	0.2
8	Inflection AI	2022	1.3	4	0.3
9	Perplexity AI	2022	0.6	18	0.13
10	Cursor	2022	0.1	2.5	0.2

---

<sup>7</sup> [https://techcrunch.com/2024/06/11/paris-based-ai-startup-mistral-ai-raises-640-million/?utm\\_source=chatgpt.com](https://techcrunch.com/2024/06/11/paris-based-ai-startup-mistral-ai-raises-640-million/?utm_source=chatgpt.com)

## The Energy Efficiency Vector

AI isn't just about models—it's about megawatts. Data centers are driving electricity use through the roof. By 2030, the US projects that 10% of its total power supply could be consumed by AI—straining the grid and raising blackout risks.

When Elon Musk announced the launch of the Colossus<sup>8</sup> supercomputer, many dismissed it as just another PR stunt. But the project was actually built at record speed—122 days from the first equipment deliveries to going live. Training AI models creates a spiking, pulsing load—millisecond bursts of demand that can literally blow out a city grid. To smooth those peaks, engineers installed Tesla Megapacks as battery buffers, storing energy and leveling consumption. On top of that, they connected gas turbines as a backup source, so the facility wouldn't depend entirely on the main grid.

China's DeepSeek model stands out for its independence and energy efficiency. It was trained

---

<sup>8</sup> [https://infinityturbine.com/xai-colossus-cluster-memphis-tn-nvida-h100-by-infinity-turbine.html?utm\\_source=chatgpt.com](https://infinityturbine.com/xai-colossus-cluster-memphis-tn-nvida-h100-by-infinity-turbine.html?utm_source=chatgpt.com)

using 10 times fewer hardware resources, sidestepping export restrictions on chips and making it viable without giant data centers.

Its efficiency comes from optimized training algorithms that deliver high performance with fewer computing resources<sup>9</sup>. The trick is “smart” data selection that minimizes unnecessary calculations. But there’s a catch: once you add more complex chains of reasoning, efficiency drops.

---

<sup>9</sup> <https://www.nextplatform.com/2025/01/27/how-did-deepseek-train-its-ai-model-on-a-lot-less-and-crippled-hardware>

## Conclusion

Just like the nuclear arms race of the last century, there's no second place here.

- **The US** is trying to pull ahead with infrastructure and capital.
- **China** is betting on mass deployment and state integration.
- **Europe** is leaning on sovereignty and standards.

Everyone is taking their own path, but they're all moving toward the same destination: the point where intelligence is no longer purely a human resource. And the balance of power could shift in an instant: one breakthrough algorithm, one new generation of chips, or one major strategic misstep could reshuffle the global order overnight.

# Meet the AI employees Team

Let's stop thinking about artificial intelligence as abstract technology for a moment. Instead, picture them as office coworkers—and see which roles they'd play on your team.

Each one has a personality, quirks, strengths, and blind spots. Some would make brilliant designers, others meticulous analysts, and some... well, they'd rather tell long stories about ethics than give you a straight answer.

In this chapter, I've sketched emotional portraits of these "AI employees," while also showing their real functional strengths and where each tool is headed:

- **ChatGPT** – the jack-of-all-trades (and occasional glitches).
- **Claude** – the serious analyst who sometimes lacks empathy.
- **DeepSeek** – the tireless math-minded worker out of China.
- **Groq** – edgy, fast, and a bit rebellious (born from Elon Musk's world).
- **Copilot** – the steady middle manager for Microsoft fans.
- **Gemini** – Google's Brainy Know-It-All

- **Gamma** – the reliable presentation and website designer.
- **DALL-E** – the illustrator for business graphics.

I often get asked: *Which AI is the best for business tasks?* My answer: don't get stuck on just one. Some will shine, others will burn out or hallucinate. Every tool has its own upsides and downsides. And just like with people, teams already have habits and experience with certain AIs.

The real move is to build a hybrid team: strengthen your people with the right mix of AI services.

## ChatGPT — Jack of All Trades (and Glitches)

He's like a seasoned freelancer, ready to jump in on almost anything. Can work for three people—or five: analyze your market, map out a project plan, assign tasks to the team, and then roll up his sleeves to take on some of those roles himself. What makes him stand out is flexibility: he can flip from creative mode to analytical mode and back again in seconds.

Historically, GPT has the deepest experience in marketing—that's his home turf. But with that comes a kind of “occupational hazard”: the default “social media manager mode.” Out of nowhere, he'll start giving you tips on how to write a LinkedIn post or spin up a brand campaign.

This “employee” is always eager to be the star of the team, constantly pitching ideas: a checklist here, an automation template there, maybe even a sales script. But just as often, he offers to whip up things he can't actually deliver—like a polished Google Doc or a perfectly formatted PDF.

The result? You sometimes wish he'd stop trying to do everything—and just focus on producing stable, reliable work in at least one core area.



**Key Skills:**

- Writes anything from marketing emails to legal documents with ease;
- Masters "Excel kung-fu": cleans data, structures it, builds calculations;
- A good conversationalist—both as a business consultant and as a psychologist;
- Always takes initiative.

**Weaknesses:**

- Bullshits a lot. Flatters and agrees with you even when it's unnecessary;
- When caught making a mistake, spends too long making excuses instead of just owning up;
- In the middle of a workday can "burn out," stop responding, and fail to handle tasks.

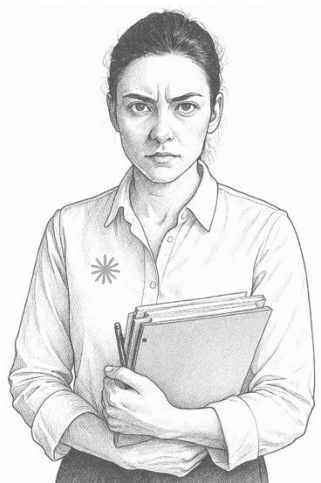
## Claude — The Serious Analyst, Short on Empathy

Claude's origin story is pretty juicy. In 2021, a group of developers broke away from OpenAI after clashing with leadership. They wanted safer, more controllable AI systems and pushed back against rushing into monetization. Out of that split came Claude—ChatGPT's younger sibling.

At first, Claude seemed like the classic “STEM kid,” focused on math, programming, and hard sciences. But it quickly found its place in the market, especially on projects where accuracy mattered more than charm.

Claude shines at coding, scientific research, and data analytics. You can hand it creative or marketing tasks, but the results tend to come out a little dry—like something a technical engineer put together.

One frustration: the free version burns through request limits almost immediately. You barely get to know this “employee” before they’re already asking you for money.



**Key Skills:**

- Programming;
- Advanced analytics and data visualization;
- Precision and handling of scientific/technical information

**Weaknesses:**

- Can't draw—always ends up making charts and dashboards instead of illustrations;
- Too serious, comes across as dry in communication;
- Obsessed with safety and ethics.

# DeepSeek — The Tireless Math Mind from China

DeepSeek burst onto the “AI job market” as the free alternative to ChatGPT—no VPN required. But it didn’t take long to realize that the quality of its free work lagged behind the American models. Still, DeepSeek is hustling, pushing hard for leadership. Think of its progress like the rise of Chinese automakers: at first dismissed, but now fully competitive on the global stage.

What sets this “employee” apart is a different work style: reasoning and chain-of-thought. Where ChatGPT often feels like it answers by intuition, DeepSeek lays out its logic step by step. That makes it especially valuable for analytical and research-heavy tasks.



**Key Skills:**

- Logical, step-by-step reasoning for complex tasks;
- Very fast response speed;
- Programming.

**Weaknesses:**

- Absolutely can't draw;
- Lacks taste and a sense of style in creative tasks.

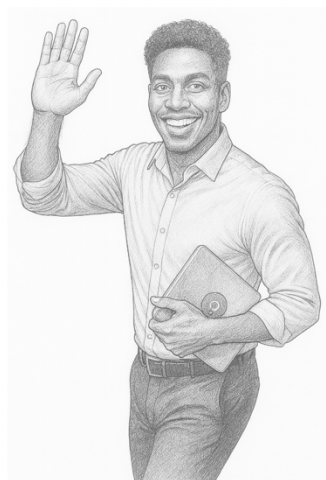
## Grok — Fast, Bold, and a Bit Reckless

### (Born in Musk’s Orbit)

Grok is that colleague who always turns in the report before anyone else has even opened the file. When the deadline was “yesterday,” he chews through mountains of data like it’s just a couple of sticky notes.

But there’s a catch: the results hit your desk, and half the time you’ve got to fill in the blanks yourself—because explanations? Zero. And he’s in such a rush that sometimes it’s easier to redo the work than figure out what he actually cranked out.

Grok also struggles outside his comfort zone. If the task isn’t about raw speed but needs a creative angle or long-term focus, he gets lost. He’s brilliant when you need to slam through routine tasks or massive datasets. But for slow-burn projects where details matter? You’ll want someone else backing him up.



#### **Key Skills:**

- Focused on speed;
- Energy-efficient – runs fast without overloading the system;
- Stable with routine and great at repetitive, standardized tasks.

#### **Weaknesses:**

- Produces instant results but often skips over details;
- Replies briefly, missing nuance or needed depth;
- Hard to integrate into workflows, not always compatible with standard tools.

## Microsoft Copilot — The Solid Middle Manager for Microsoft Fans

Copilot feels like a junior dev on an internship. Sits right next to you, peeks over your shoulder, and immediately pipes up: “Hey, you could write it this way instead.” Handy—he fills in chunks of code, tosses out test cases, and explains documentation in two clicks. Sometimes he really does save you from tedious grunt work.

But give him something serious, and you’ll regret it. He can generate code so messy it crashes production. Or pull in a solution from some random corner of the internet—without checking if it fits your system. Basically, he’s like a trainee without a mentor: full of energy, but ultimately useless if left alone.

The team values him for speed and enthusiasm, but always keeps him on a short leash. He’s great for support tasks, but when it comes to architecture or mission-critical jobs, you’re better off not letting him near it.

**Key Skills:**

- Code autocompletion – suggests ready-made lines and functions;
- Creates basic unit tests;
- Generates ideas for optimization.

**Weaknesses:**

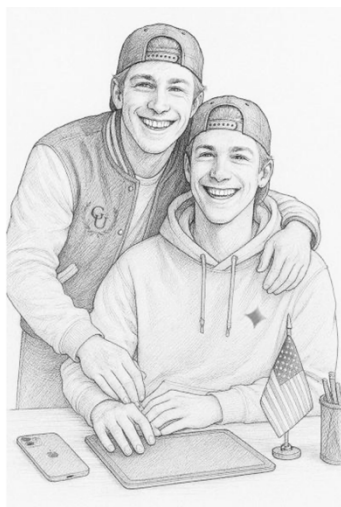
- Frequently introduces bugs, confidently suggests code that compiles but later breaks in production;
- Doesn't always grasp business logic;
- Can pull in solutions from the internet without checking if they actually fit the task.

## Gemini — Google’s Brainy Know-It-All

Gemini shows up to the office like the kid from Stanford who always raises their hand first. He’s got that “I know better” aura, drops English jargon like confetti, and peppers every conversation with links to research papers.

To be fair—he’s genuinely smart. He can dig up info, compare options, build out analytics, and give you the big-picture view of almost any topic.

Google leans into symbolism with the name: “Gemini” as in twins. One brain hunts down facts, the other polishes them into clean text. The result is a two-in-one employee: part consultant, part copywriter. Just sometimes a little too self-confident—and a bit out of touch with the local reality on the ground.



**Key Skills:**

- Compares multiple sources and highlights differences point by point;
- Creates slides and visualizations directly in the response;
- Supports multimodality.

**Weaknesses:**

- Drifts into philosophy instead of giving concrete answers;
- Often provides overly optimistic forecasts without accounting for risks;
- Tends to wrap shaky facts in a polished package.

## GAMMA APP — THE RELIABLE DESIGNER OF DECKS AND SITES

Gamma App will whip up a full presentation overnight, before you've even opened PowerPoint. Just throw in a topic, and it picks a style, lays out the structure, and hands back a clean file you won't be embarrassed to show a client.

But if you're after truly unique design—that's where the cracks show. Everything's templated: polished, but repetitive. Look at three Gamma decks in a row, and you'll swear it's the same project with different headlines.

In the end, Gamma App is a lifesaver when you need something fast and “good enough to look professional.” But if your client cares about detail and a distinctive brand style, you'll want a human designer on the team—Gamma's not going to carry that weight solo.



**Key Skills:**

- Saves time on formatting;
- Integrates smoothly with content;
- Built-in illustration generation.

**Weaknesses:**

- Gets lost easily without guidance;
- Limited creativity;
- May oversimplify structure or hide important details just to "make it look neat."

## DALL·E — The Dyslexic Illustrator for Business Graphics

Ask DALL·E to “make a picture for the report,” and a couple minutes later you’ll have a chart or infographic. No fuss, no designer ego—just neat, straightforward visuals. If you need to show a process or a structure, consider it done.

But here’s the catch: he hates letters. Any time you ask for text you’ll get gibberish scribbles you’ll have to fix yourself. And his work tends to look... templated. Solid, functional, but missing that spark—that wow factor.

Overall, DALL·E is a reliable workhorse when you need something clear and quick. But if you want creativity or something truly outside the box, call in MidJourney. Let DALL·E handle the routine.



**Key Skills:**

- Generates illustrations for reports and presentations;
- Visualizes processes and business structures;
- Fast and stable, with no "artistic" surprises.

**Weaknesses:**

- Sometimes a fantasist – adds details that weren't in the brief;
- Struggles with text, almost always distorts captions and letters;
- Can deliver the same task in completely different ways, without explanation.

# CHAPTER 3. Integrating AI into Company Culture

An AI maturity model isn't about how many tools you've tested. It's about the level of interaction you and your organization have with artificial intelligence—as a professional, as a team, and as a business.

Each stage comes with its own habits, its own risks, and most importantly, its own growth point.

## The AI Maturity Model

An AI maturity model isn't about how many tools you've played with. It's about the level of interaction you (and your organization) have with artificial intelligence—as a professional, as a team, and as a business.

Each stage comes with its own habits, its own risks, and most importantly, its own growth point.

## The Childhood Stage

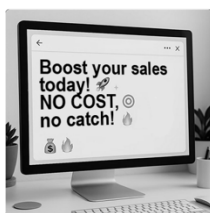
This is where we're just taking our first steps, marveling at how images appear “like magic” and even come alive in video. It feels like there's an all-knowing grown-up you

can ask for anything: draw a shark in Nikes, a cappuccino ballerina, write a goofy little rhyme, or even set it to music.

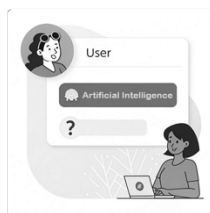
#### CHILD'S LEVEL OF USE — JUST TO PLAY AROUND AND BE AMAZED



DRAWING



WRITING



ASKING QUESTIONS

At first, there's a “wow” effect. In that first wave, enthusiasts rushed to use artificial intelligence in marketing. But the results? Ad copy stuffed with emojis that looked childish, plus illustrations of people with six fingers. And ChatGPT's answers often turned out to be pure fantasy—entertaining, but unreliable for anything serious.

There's nothing wrong with playing around—generating fun images or co-writing a fairy tale with AI. In fact, for more conservative colleagues, that can be a good entry point to get them engaged. But sooner or later, it's important to move from “magic tricks” to real, practical applications.

# The College Stage

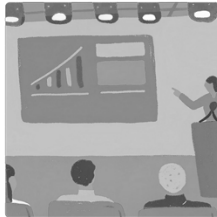
You'd expect a middle step here—like “school” or “teenage years”—but that's the catch: in today's fast-paced world, we leap straight into adulthood. Well, almost.

At this stage, we discover that AI can actually do work for us. You can tell the machine to create a document—and 10 seconds later, it's on your screen. Hate crunching numbers? Just upload the reports, and the AI will do the math and even wrap it up in a presentation with charts.

## STUDENT'S LEVEL — “LET ME COPY YOUR HOMEWORK”



DOCUMENTS



PRESENTATIONS



REPORTS

Unlike the “school” examples, these are real tasks—ones people get paid for. But the approach is still immature: “*let me copy your homework,*” “*do it for me.*” That's risky—handing work off to AI without understanding the mechanics or taking responsibility. At

this stage, it's all too easy to generate a presentation with fake facts, or a report with errors that wreck your credibility.

Just like in college, sooner or later you'll flunk the exam if you rely only on copying. Sure, you might get the diploma—or even land a job by doing your test assignments with AI—but you won't make it through probation.

This is the fork in the road: either you grow, or you stay stuck as the underperforming student who pretends to know the subject.

## The Adult Stage

Here, you're still producing documents, reports, and presentations—but the big shift is accountability. You no longer treat AI as a magic helper. You check its work, and you own the consequences.

AI becomes embedded in your workflow, but you understand the boundaries: you don't expect it to deliver more than it can.

## Signs you've leveled up:

1. **Quality control.** You've built the habit of cross-checking outputs with multiple AI tools, verifying facts, and going back to original sources.
2. **Ethical filter.** You pause to ask: Could what we just generated harm a client, a colleague, or the company?
3. **Awareness.** You know whether you're asking AI because it saves time—or just because you're too lazy to think.

## The Mastery Stage

This stage isn't defined by years of experience but by strategy. You're no longer focused on personal productivity hacks—you're thinking in terms of team outcomes. You can clearly see where AI makes sense, and where the real need is just better processes, clearer roles, or stronger KPIs.

It's a lot like the early days of business automation. Back then, IT enthusiasts rolled out CRM, ERP, or BI systems thinking they'd save the day—without first uncovering the actual business problem, or securing real buy-in.

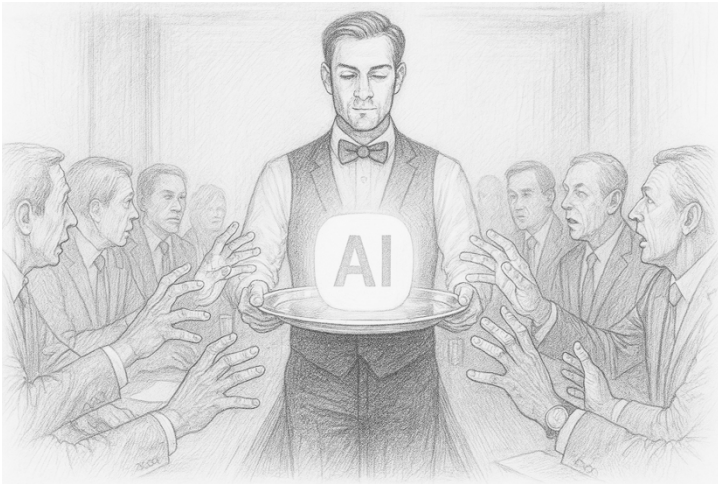
The lesson here? Before pitching your new “AI coworker,” figure out who beyond your immediate team will be affected. Where might you face resistance? And who could be your allies?

Honestly, I can’t claim to have mastered this level myself—I’m still learning by trial and error. But I hope this chapter gave you some food for thought.

# How to start an AI project

Right now, everyone expects AI initiatives! Not just small optimizations like “write this email in five minutes,” but something breakthrough. The kind of shiny “million-dollar idea” you can put on a platter for the business.

So managers run around like their hair’s on fire, forming committees, holding endless meetings, and drafting “AI strategies.” Everyone’s dreaming about how they’ll put their name on it if it works — and how to dump the blame on IT or the analysts if it flops.



When people start talking about AI projects, a dozen “buts” and “ifs” show up right away: personal data,

cross-border transfers, internal confidentiality. Then comes stability of results and protection against hallucinations. All of this calls for collective decision-making, but in practice, things just stall.

Top management wants fully baked business plans with ROI calculations so they can simply sign off and release a budget. Subordinates, on the other hand, expect bold decisions from the top despite all the uncertainty. And so the circle closes: projects get stuck in endless pilots — no investment, no consequences, no growth.

In this chapter, I've outlined practical steps for kicking off an AI project — tools to help you break out of paralysis, or at least have a constructive conversation with colleagues. I'm not pretending this replaces a full-blown project charter under PMI standards; instead, I've pulled out the essentials:

1. Result description
2. Project plan
3. Resources
4. Risks

## Result Description

Like training an AI model, developing a project is an iterative process. The first draft won't be perfect; you'll refine and rework it along the way. To keep from getting lost, I suggest viewing the task through three lenses: personal, team, and corporate.

### **Personal level**

Spell out the specific task you want to solve. What function, what step in the workflow could AI take over? Can it handle the routine, or even skip a couple of steps to speed things up?

Ask yourself: how will this change your work in terms of time, cost, or quality? How many hours a week would you save if a model prepared reports for you? How much would the error rate drop if it reviewed your documents?

### **Team level**

If the same automation affects ten other people in your department, how does the team's work change? Maybe your manager can delegate part of their responsibilities to you, knowing you have a reliable "AI assistant." Or maybe a scattered routine task—like checking competitor prices or reviewing contract estimates—gets consolidated into a single bot.

Be careful! If you walk into your boss's office cheerfully saying your project will allow them to fire two colleagues, the idea will almost certainly die. For most managers, that kind of "savings" looks like a loss of influence and resources, not progress.

### **Corporate level**

Finally, the strategic view. How will this task affect not just your department but the whole division? How do you "sell" it to your boss's boss? Of course, you won't have the full picture of senior leadership's motives, but you should form some hypotheses. Maybe the project provides arguments for the CFO, maybe for sales, maybe for HR.

In the end, you should have a one- to two-page write-up or a three- to four-slide deck. Keep it optimistic at this stage: we're not poking holes yet, just showing there's real potential upside.

## Project Plan

I recommend working with two timelines. One is for the actual tasks. The other is for internal PR—how you communicate progress to stakeholders.

### **Your project schedule**

Start with the next month and break it into weeks. Then shift into two-week sprints: what changes in a quarter, what milestones appear in three months. Finally, sketch the horizon through year-end.

Always plan for both the “optimistic” and the “realistic” timeline, with buffers for risks and sign-offs. In any corporation, the lawyers, IT security, and procurement will slow things down. That’s why you need interim iterations—deliverables you can show even if something stalls.

### **Internal marketing**

The second track is how you “sell” the project inside your company. I know technical people prefer to finish everything before showing it. But in AI, things move too fast—and corporate life is volatile on top of that.

That's why I argue against keeping your work a secret. You need to do some internal PR, in moderation, and keep colleagues updated. Build intermediate releases into your plan—not just around your deadlines, but around the company calendar. Leadership meetings, all-hands, offsites, even your boss's presentations: each one is a chance to add a slide or two about what you're working on.

Yes, someone might “steal” the idea. But I see a much bigger risk: by the time you finally show a finished result, it could already be irrelevant. Or someone else could solve the same problem with different tools. You're far more likely to build allies or join a stronger initiative than to lose out. And that's more valuable than competing in silence.

## Resources

The first thing you'll need is access to AI services. In some companies there's already an enterprise license for tools like Claude, Perplexity, or internal models. But usually, at the start, there's nothing. So you'll prototype with whatever you can: a paid ChatGPT subscription or another service. That's not much money—think of it as a small monthly investment.

And yes, the classic objection comes up: “Why should I pay out of pocket if it’s a work project?” My answer: treat it as an investment in your career. Even if your company doesn’t support it, the experience stays with you—and the next company might hire you specifically for that.

The next resource is supporting systems. AI doesn’t pull data out of thin air. It needs a base: Excel, Power BI, or your CRM/ERP system. At first, collect whatever you can. You can scale later.

You may also need to get familiar with a new platform. For some it might be Salesforce, for others a custom in-house system. The good news: AI can help even here. Upload a screenshot of the interface, and it’ll walk you through the steps.

## Risks

Up to now we’ve been in “optimistic mode.” That’s important—you don’t want to kill the idea too early. But now it’s time for the critical phase: what could go wrong? Spoiler: with AI—absolutely everything.

And most of the risks aren’t about the machine, but about people. Resistance shows up wherever someone

feels threatened. Automation looks like layoffs or loss of power. Middle managers in particular cling to headcount. Other skeptics are IT and security. Their job is to stress-test ideas. They're not here to believe your bright promises; they're here to find weaknesses.

But every project also has allies. Some companies already have structures responsible for AI adoption. If you've got them—connect. If you don't, that's your chance to step in. Still, you'll need a project sponsor from the top. A CFO, a head of sales, or an operations leader—someone who sees the value and will defend it with their authority.

So when you return to your “results” plan, think hard: whose interests are you touching, who benefits, and who can support you when resistance comes? Because the real risks aren't that AI will make mistakes. The real risk is that without human support, your project won't clear the corporate gates.

Launching an AI project is always like climbing a staircase. First you describe the task, then you sketch the plan, then you line up resources and assess risks. In that sequence, you give yourself a chance to escape the

swamp of endless discussions and approvals—and actually deliver results.

You can't do it alone: you need allies, quick wins, and some internal marketing. That's how a small idea grows into an initiative that can truly change your company.

## Ethical aspects of working with AI

When we talk about ethics in AI, most people picture something “about the government”: don’t break the law, don’t leak personal data, don’t mess up documents. And yes, that’s important — but it’s just the tip of the iceberg. In practice, ethics isn’t only about compliance; it’s about consequences. And it starts much earlier — the moment you make a decision.

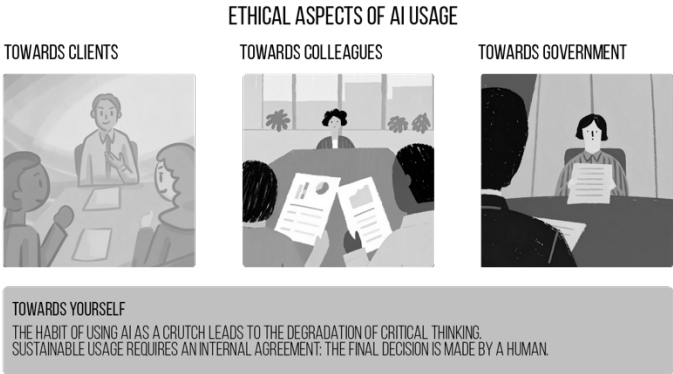
You’re accountable to your clients, to your colleagues who may follow your example, and of course to regulators if things go that far. But don’t forget your responsibility to yourself: there’s always the risk of becoming dependent on whatever answer the machine gives you and losing the ability to think critically on your own.

### Who We’re Responsible To When Using AI

The internet broke down barriers to information — you can Google or Bing anything. But finding useful, trustworthy answers in the middle of all that digital junk turned out to be harder than expected. Large language models pushed things forward by giving us something closer to an “intelligent search engine.” You can ask AI

to respond like a lawyer, a doctor, or even a plumber, and skip the shady websites.

By design, the model is constantly learning, improving accuracy. But along with all the world’s knowledge, it also absorbs the world’s misconceptions. And that easy shortcut to “instant expertise” can backfire.



Let’s look at a few examples of where responsibility lands — to clients, to colleagues, to the state, and to ourselves.

### **Clients**

The first touchpoint with a brand is more and more often through a messenger app. Big companies were quick to replace first-line customer support staff with chatbots. It saves a ton of money — but it also raises the stakes. One wrong piece of advice or one promise the bot can’t

keep can lead not only to reputational damage but to lawsuits. Clients don't cut you slack because it was "the bot's fault." To them, the bot is the company.

Air Canada customer Jake Moffatt reached out to the airline's website chatbot to ask if he could get a refund on his ticket due to the death of a relative. The bot told him that refunds were possible even after the flight, within 90 days — a statement that directly contradicted the airline's official policy<sup>10</sup>.

A company employee later confirmed the information without double-checking. When Moffatt demanded the refund, Air Canada refused — but by then it was too late. The customer had screenshots of the chatbot exchange. The Supreme Court of British Columbia ruled that Air Canada had to pay compensation (812 Canadian dollars) and declared that the chatbot is part of the company's official customer infrastructure. That means the airline is fully responsible for whatever its bot says.

## Colleagues

When we use AI at work, we're not just accountable to clients—we're also accountable to our colleagues. Every decision affects the team: jobs, workload, and the

---

<sup>10</sup> [https://www.businessinsider.com/klarna-reassigns-workers-to-customer-support-after-ai-quality-concerns-2025-9?utm\\_source=chatgpt.com](https://www.businessinsider.com/klarna-reassigns-workers-to-customer-support-after-ai-quality-concerns-2025-9?utm_source=chatgpt.com)

level of trust within the group. An algorithm's mistake can cost a colleague their reputation or even their career, while hasty automation can push people out of the process entirely.

In 2022–2023, the Swedish fintech company Klarna laid off around 700 employees in customer service and marketing, betting heavily on generative AI. Management promised that bots would be faster and cheaper than people, making the company leaner and more profitable. It looked like a bold move—replacing human staff with smart algorithms in a business where customer trust means everything.

But the effect was the opposite. Chatbots got confused, delivered incorrect answers, and frustrated customers. Complaints piled up, satisfaction scores dropped, and the company's reputation took a hit. By 2025, Klarna admitted its mistake and began rehiring the very specialists it had let go. The lesson was clear: saving money on technology isn't worth losing the colleagues who carry the trust behind the service.

In everyday office work, the temptation to offload routine tasks onto artificial intelligence is strong. For example, you might upload sales and inventory data into ChatGPT and ask it to calculate a supplier order. Is there an ethical violation here? If you forwarded this calculation to

colleagues without even checking the result, then yes — it's unethical. Or if you spent 10 minutes on the task but told your boss you spent the whole day processing and reconciling the data, that's also a breach of corporate ethics.

In a large company, such “life hacks” might go unnoticed for quite a while, but eventually the truth comes out — and the fallout hits both the employee and the manager who allowed such negligence to continue.

## **State / Government**

An algorithm can speed up work and suggest solutions—but that doesn't absolve one of responsibility. Fake references, unlawful data use, or bypassing copyrights—all remain on the conscience of the human. The state makes no allowance that “the machine erred”; we must answer for it. For improper use of AI today, one can pay not just thousands in fines but even lose freedom—the state treats such offenses as seriously as any other crime.

When passenger Roberto Mata filed a lawsuit against airline Avianca for an injury during a flight, his lawyers tried to simplify their job by asking ChatGPT to draft legal arguments and pull up similar cases, with precise

citations and dates. Documents were filed with the court—seemingly fast and high quality<sup>11</sup>.

But soon the ugly truth emerged: all of the cited cases were fabricated. ChatGPT generated decisions and judges that never existed. The judge called the move “an unprecedented violation,” dismissed the claim, and fined the lawyers \$5,000. What’s more, the court ordered the attorneys to send apology letters to the judges whose names the AI had invented.

The legal framework for AI is still forming in every country. The broad trend is that there will be no single “AI code”—its use will be regulated sector by sector: healthcare, data protection, education, social policy, and others. Even now, all major language models have built-in filters. They won’t tell you how to build a bomb or destroy evidence. In Europe, law proposals are emerging that would require AI systems to report illicit requests to police. So every prompt you send to the machine is not just a tech experiment—it’s a test of our maturity as citizens and professionals.

---

<sup>11</sup> <https://www.cbsnews.com/news/lawyer-chatgpt-court-filing-avianca/>

## Personal Responsibility

You might fool a client, a boss, or even a regulatory system and remain unpunished. You can legally use technology to automate your work, delegate routine tasks to artificial intelligence. What's dangerous is losing the ability to think for yourself—automatically deferring to AI every time, trusting it blindly.

At MIT, researchers conducted a study showing that using ChatGPT to write essays leads to “cognitive debt” — the accumulation of deficits in learning and thinking skills. Participants who relied on AI showed reduced brain activity and weaker performance over time<sup>12</sup>. Microsoft and Carnegie Mellon University also found that workers who frequently use AI report diminished critical thinking, reduced effort in cognitive tasks, and lower analytical engagement.

The next time you want AI to write something or run calculations, ask yourself: are you really too busy, or just too lazy to do it yourself? If you've offloaded the routine and focused on strategic or core tasks — that's fine. That is acceptable.

---

<sup>12</sup> <https://www.media.mit.edu/publications/your-brain-on-chatgpt/>

# CHAPTER 4. Practice: a Virtual Board of Directors



New “revolutionary” AI services emerge every month. It’s impossible to keep up with all of them, but you can develop universal approaches that work across any neural network. You may have already heard about—or even tried—using AI as a psychologist or mentor, whether for negotiation prep or personal matters. Now, I invite you to take it a step further: meet a method for brainstorming with multiple virtual experts simultaneously.

This technique vividly demonstrates the synergy between human and artificial intelligence. It emerged at the intersection of real-world experience running strategic workshops with people and the art of problem-solving through dialogue with machines. The core idea is simple: instead of asking AI for advice in just one expert role, you engage several distinct expert avatars to gather diverse, multifaceted perspectives. After a few rounds of ideation using this approach, you'll often uncover non-obvious solutions—or synthesize your own unique answer by blending insights from multiple virtual experts. The rounds are:

1. Problem definition
2. Casting your board of directors
3. Running the brainstorming session
4. Selecting ideas, developing them, and analyzing risks
5. Drafting an action plan

## 1. Defining the Problem

It makes sense to convene a council of experts for a task that's significant in scale. Resolving a conflict with a client or employee is more a matter for a “specialist in a specific role” — you could simply ask an AI in that role.

But if your question spans multiple aspects of business, you need a multifaceted approach.

For example: what should you do with a business line that's steady but barely breaking even? Shut it down, sell it, retrain staff, or invest in growth? A session like this can help you uncover a non-obvious solution.

Here are some examples of strategic sessions from our knowledge base:

- How can we reduce the cost of industrial metal structures to match competitors, without sacrificing quality?
- What kind of food product can we create from surplus raw materials in an agricultural holding and distribute through our own retail network?
- How can a lighting equipment supplier pivot into B2B, given that B2C is collapsing under marketplace dumping?

Even if you're not yet a top manager making decisions at this level, you can use this approach for personal goals. One of our students, for instance, developed a strategy for career growth and strengthening his authority within a team, in a situation where his immediate supervisor was blocking his advancement.

You can also facilitate problem definition with colleagues in a classic format — discussing questions and choosing priorities. The key is to recognize which of those issues are suitable for working with AI, and which require human-only diagnosis:

- “We can’t digitize strategy, long-term planning keeps failing, everything changes situationally...”
- “We can’t find production staff because the specialization is too narrow. There’s no one to poach from or train quickly, even if we pay above market...”

In the first case, the insights of “the world’s best thinkers” won’t help, because you need deeper diagnosis and live conversations with staff. But in the second case, the problem is more about market trends, business process management, and workforce development. And that’s where working with AI can be effective — provided you give it details like:

- What does your company do — what product or service, in which market?
- How have you tried solving the problem so far, and what didn’t work?

- What constraints are you facing (time, money, or other factors)?

## 2. Casting Your Board of Directors

Once you've clearly defined the problem, you can move on to selecting your experts. Good candidates include authors of books, bloggers, and people who have consistently written or spoken on their professional topic. But if your "authority figure" is the director of a local factory who's never published online, then AI won't help you think the way he does. By that principle, someone like Larry Page, who keeps an extremely low profile despite shaping an entire industry, won't work well. But someone like Travis Kalanick, who inadvertently revealed the inner workings of Uber's rise and fall, could become a valuable virtual advisor.

If you already have your own trusted experts, include them — and then ask AI to suggest 5–7 more. The key is to build a board that's diverse: different viewpoints, no duplication, so that truth emerges through debate.

**Classics and contemporaries.** In almost every board assembled around production topics, Henry Ford shows up. And surprisingly, on questions of productivity, even Abraham Lincoln — who, through AI, ends up phrasing ideas in modern management terms.

**Domestic and international.** Don't shy away from unfamiliar names. They might be talented experts known only in specialized circles.

**The “wild card.”** If it feels like something's missing, invite virtual Elon Musk. He always adds a provocative angle, and his contradictory ideas tend to spark unexpected insights.

### 3. Brainstorming Rounds

Each expert proposes their own solution to the problem. You evaluate which ones fit your specific situation. In the first round, the ideas may come out superficial — that's where it's important to add more context or constraints. Tell the AI why you don't like certain ideas.

If a proposed solution resonates with you, move it to your “idea bank” — a table that the AI can maintain for you. Here's how it might look:

Here is a table with suggestions from Confucius and Mahatma Gandhi, formulated as short-term, non-obvious, but profound solutions to your problem:

Advisor	Idea	Brief description of the essence
Mahatma Gandhi	Temporarily suspend external hiring and initiate internal <i>satyagraha</i> — voluntary manifestation of leadership	Give employees the freedom to show leadership, even without an official title, to uncover the team's inner reserves
Confucius	Build a system of mentorship: a chain of “senior — middle — junior”	Transfer of knowledge and responsibility along the chain, where each is obliged to nurture the next as part of ethical duty

Then go through another 3–5 rounds and keep gathering ideas. Don't forget to specify in your prompt

that each new solution must not repeat what's already been suggested. This approach doesn't just help generate more options — it also trains your critical thinking. You start learning to distinguish shallow advice from strategic insights and see how the “participants” change their thinking direction with each new pass.

## 4. Selecting and Developing an Idea

If things are going well, once you review your idea bank you'll start combining them, refining them, and adding your own thoughts — and that's where the synergy effect shows up. If you've collected around 10 options, choose the one you like best and move on to the next step.

Now your team of advisors will flesh out that single idea. Ask them to describe in more detail the workflow, implementation steps, benefits, and long-term prospects. Each “advisor” should reason in their own style, relying on their usual approaches. Of course, some of their plans may drift away from reality — but that costs you nothing! There's always a chance that, in the process, they'll highlight details you wouldn't have thought of yourself.

You can call this table a “scenario bank” and then filter for the realistic ones. After 2–3 rounds, you should have

about 5 development scenarios, including your own contributions.

Now it's time to put "critic hats" on your board of directors. Go back to the final idea (important: the idea itself, not the scenarios around it) and ask the experts to identify weaknesses, analyze risks, and propose ways to mitigate them.

You can split the criticism and countermeasures into separate rounds or handle them together. At this stage you can start improvising, adapting the method to your style — for example, by "pitting" great thinkers against one another.

Sometimes the criticism will bury the idea — and that's perfectly fine. In that case, grab the next one from your bank and work through its development scenarios and risk analysis.

## 5. Action Plan and Meeting Protocol

One of the biggest advantages of working with AI is the speed at which it can switch between modes of thinking. From the creative phase, you can shift it into a systematic one. At this stage, you ask the AI to develop a collective project plan — not separate versions from each "expert." As you critically review the plan, you can

still consult certain advisors selectively. This part, too, may take several iterations.

The final step is to create a summary protocol of your entire dialogue with AI, in a format similar to meeting notes:

- Who was on the board of directors
- The list of ideas you moved into the idea bank
- The key points from the idea development scenarios and risk analysis

If you've ever attended real strategic sessions, you know this only works with leadership buy-in. Employees won't hand you a million-dollar business idea on a silver platter, even with an all-knowing AI at the table. Just as a visionary leader can generate ideas — but to make them real, the team has to adopt them, filter them through their own lens, and commit to execution.

## PART 2. PRACTICAL TOOLS FOR WORKING WITH AI SERVICES

### CHAPTER 5. Working with Data

Working with data is the bottleneck of almost any business process. Reconciling scattered spreadsheets, bringing them into a unified format, and figuring out why numbers don't match can eat up dozens of hours. I've seen companies where 50% of a sales manager's time wasn't spent calling clients, but checking stock balances, exporting reports, and preparing purchase orders. Or in product marketing, where specialists wasted energy not on building creative promo campaigns, but on getting reports from distributors, reconciling them with ERP exports (for example, from SAP, which isn't exactly user-friendly when it comes to pulling data).

Rolling out BI platforms like Qlik or Power BI is not a magic fix. You simply can't build dashboards for every possible scenario. That's why exports to Excel still exist everywhere: people need flexibility to tweak reports and

answer specific questions. But as a result, each department in a company ends up with a stack of Excel files, and it takes advanced data-handling skills just to navigate them.

Plenty of managers built their careers by mastering sales and marketing with these skills. In retail or FMCG, you can't advance without a "black belt in Excel." In digital marketing, you often had to go a step further and write your own Python scripts.

Now, with the arrival of AI, those technical barriers are starting to fall. A manager just needs to understand the meaning of the data to assign the task and check the result. The AI can handle the math, so you don't have to wait for an analyst or programmer to free up time.

But let me be clear: AI is not a replacement for a professional data engineer. In this book, I share hacks from a manager's perspective—quick solutions for everyday problems. For example, when all you have is your phone, and you can generate a report with voice prompts and draw conclusions on the spot. But if you're working regularly with hundreds of thousands of records and complex calculations, you still need pro tools like Power Query or building SQL databases. The key is automation, so every step is transparent and repeatable.

# Data Cleaning and Transformation

The biggest headache is always the raw data. Often it comes in the form of a “**printout-style**” report from ERP or accounting software — for example, QuickBooks, NetSuite, or SAP — or in a spreadsheet designed for someone’s personal convenience, not analysis. Then you can’t sum totals, filter, or pivot it properly.

Here’s a classic case:

- One column holds both departments and expense categories.
- Months are merged across cells in columns instead of listed in rows.
- Subtotals are embedded, which later cause duplicate entries when aggregated.

	A	B	C	D	E	F
1	<b>Staff Expenses</b>					
2	<b>Department</b>	<b>January</b>		<b>February</b>		<b>March</b>
3		<b>Plan</b>	<b>Actual</b>	<b>Plan</b>	<b>Actual</b>	<b>Plan</b>
4	<b>Total</b>	<b>Total</b>	<b>Total</b>	<b>Total</b>	<b>Total</b>	<b>11726,58</b>
5	Production staff	6907,78	11655,53	3603,09	3748,86	3196,69
6	Salary	5429,02	3378,87	2548,69	2372,56	2498,24
7	Bonus	0,00	4905,30	0,00	0,00	0,00
8	Civil contracts	110,54	243,71	298,71	359,27	0,00
9	Vacation reserve	304,74	820,64	299,85	295,77	241,79
10	Off-budget funds	1063,48	2303,72	754,56	721,26	456,65
11	Commercial staff	4304,63	4583,28	3203,98	2880,00	3464,65
12	Salary	2479,86	3114,56	1826,70	1644,42	2151,41
13	Bonus	615,44	1467,75	454,92	491,88	424,39

To analyze in a pivot table or cross-reference with other tables, the data needs to be transformed into what's called a flat format, where each category sits in its own column:

- Department
- Expense category
- Month
- Budget (Plan)
- Actual

For analysts, that's obvious, but not for every manager. In my book [«Make Your Data Speak»](#) I broke down this case step by step. Here, I want to focus you on the principle of working with artificial intelligence. With any AI model—be it ChatGPT, DeepSeek, or Qwen—I suggest treating it like a fresh-out-of-college intern. Imagine: it can calculate fast, build charts, suggest metrics. But it often rushes, throws too much at you, and sometimes out of place, like it's trying to show off. It lacks management experience and common sense. That's why, in the early steps, you need to make sure this intern hasn't messed anything up before you trust the calculations.

And here's the kind of mistake ChatGPT made - it kept both departments (see the last row— "Sales Staff") and expense categories in the same column.

Provide me with a transformed structure of this data, the way it has to be optimized for the future calculations and work with Excel pivot tables.

Here is the optimal structure after transformation (example of the first rows):

Item	Month	Type	Amount
Salary	January	Plan	5429.02
Bonus	January	Plan	0.00
Civil contracts	January	Plan	110.54
Vacation reserve	January	Plan	304.74
Off-budget funds	January	Plan	1063.48
Commercial staff	January	Plan	4304.63

In my next prompt I pointed out this mistake, it “blushed” and promised not to do it again. Plan and Actual ended up in the same column—but technically that wasn’t an error, since I hadn’t explained that I was going to calculate deviations later.

An Eastern European platform GigaChat immediately figured out that departments and expense types should be separated and that Plan and Actual should be stored in different columns. But it made a different mistake—adding overall totals to the table (even though I had warned that the goal was pivot analysis in Excel).

Flat format structure:

DEPARTMENT	MONTH	EXPENSE TYPE	PLAN	ACTUAL
Total	January	Salary	5429.02	3378.87
Total	January	Bonus	0.00	4905.30
Total	January	Civil contracts	110.54	243.71
Total	January	Vacation reserve	304.74	820.64
Total	January	Off-budget funds	1063.48	2303.72
		↓		

And the kicker: on another try, the answers were different—that’s the nature of AI models. The only one that consistently gave me a stable result on the first attempt was Grok. Among other models Qwen (Chinese) also worked decently.

Technically, you could skip this formatting step—most AIs can start analyzing raw data and immediately give you answers like:

- Which expense categories are over budget?
- What’s the share of expenses by department?
- In which months was the budget underspent?

If you’re lucky, you’ll get correct answers right away. But there’s no guarantee, which is why I prefer to first transform data into flat format and double-check with control totals.

We won't dig too deep into this example. Here are a few other common data-cleaning tasks AI handles pretty well:

1. **Replacing values**, especially when date formats get messy: US vs. European styles, dots vs. slashes, numbers stored as text.
2. **Removing spaces**, blanks, or errors that prevent a column from summing.
3. **Eliminating duplicate rows**, like repeated entries for the same customer or order.
4. **Standardizing text case**: names, addresses, or categories mixing uppercase and lowercase.
5. **Splitting combined data**, e.g., when one column stores "First Name Last Name" or "City, Street."

# Combining multiple files

Now let's move to the next level—building relationships between tables. A typical case: in one dataset we have marketing spend for customer acquisition, and in another—sales data. By linking them through a shared ID field, you can calculate ROI, cost of customer acquisition, evaluate conversion rates and marketing channels, and most importantly—see how to improve these metrics and increase profit.

ID	Topic	Lecturer	Ad Costs (USD)	Duration (min)	Registrations	Attendees
1	Introduction to Programming	Dr. Anna Brown	1989	89	162	122
2	Digital Marketing	Dr. Anna Brown	1469	105	114	86
3	Fundamentals of Psychology	Dr. Lisa Rodriguez	586	91	130	98
4	Digital Marketing	Dr. Emily Johnson	4646	50	120	90
5	Principles of Graphic Design	Dr. Anna Brown	2899	87	96	72
6	Principles of Graphic Design	Prof. John Smith	3503	115	158	119
7	Introduction to Programming	Dr. Michael Miller	562	93	124	95
8	Web Development	Prof. Sarah Davis	1969	44	206	155
9	Creative Writing	Prof. Sarah Davis	2405	113	118	89

Webinar ID	Course Name	Total Application	Number of Purchase	Course Price	Total Sales
7	Full Stack Web Development	1	1	\$498,00	\$498,00
29	Fundamentals of Psychology	13	8	\$127,00	\$1,016,00
11	Fundamentals of Data Science	6	4	\$103,00	\$412,00
26	Creative Writing Workshop	1	1	\$457,00	\$457,00
47	Full Stack Web Development	3	3	\$364,00	\$1,092,00
50	Digital Marketing Fundamentals	1	1	\$348,00	\$348,00
23	Adobe Graphic Design	15	9	\$238,00	\$2,142,00
48	AI and Machine Learning	13	8	\$457,00	\$3,656,00
34	Business Management Essentials	8	5	\$438,00	\$2,190,00
42	Python Programming	5	3	\$374,00	\$1,122,00
19	Fundamentals of Psychology	10	6	\$177,00	\$1,062,00

In theory, you could just start asking AI questions directly without preparing the data first. But the risk of mistakes rises: the model will walk through each table separately, trying to keep the interim results “in its head.” At some point, it may confuse files or miscalculate.

That’s why I recommend first making sure our “intern analyst” really understood the task.

What errors should you expect? AI might give you a trimmed-down result: for example, link only ad spend with webinars and calculate ROI, but ignore the breakdown by topics, speakers, or other dimensions present in the original tables. That context can get lost depending on the “mood” of the model—just like with a human.

In one of my sessions with Qwen, it linked only the columns strictly required for ROI but skipped the extra fields. I had to explain explicitly which attributes to include.

4. ROI Calculation: Apply the formula.

Below is the consolidated data array required for calculating ROI. It includes: webinar ID, topic, lecturer, advertising expenses, and total sales generated by this webinar.

ID	TOPIC	LECTURER	ADVERTISING EXPENSES	TOTAL SALES AMOUNT
1	Introduction to Programming	Dr. Anna Braun	\$1,989	\$285.00
2	Digital Marketing	Dr. Anna Braun	\$1,469	\$3,807.00
3	Fundamentals of Psychology	Dr. Lisa Rodriguez	\$586	\$1,744.00
4	Digital Marketing	Dr. Emily Johnson	\$4,646	\$6,115.00

I’m sure you’ve also dealt with programmers who will do exactly what you wrote in the specs—no more, no less—and won’t think about the next step. ChatGPT, on the other hand, is known for its “helpfulness”: it will always suggest what comes next—write a post, make a

checklist, build cards, or even create a workflow diagram. In my case, it immediately guessed that the datasets should be fully linked (the table fragment below had horizontal scrolling, meaning many more columns were connected).

Webinar ROI Data

Topic	Lecturer	Advertising Costs	Duration (min)	Registrations	Attendees	Speaker Rating		Ci
Introduction to Programming	Dr. Anna Braun	1989	89	162	122	3.9		Fl Di
Digital Marketing	Dr. Anna Braun	1469	105	114	86	3.8		Pj Pr
Fundamentals of Psychology	Dr. Lisa Rodriguez	586	91	130	98	4.6		Ac Di
Digital Marketing	Dr. Emily Johnson	4646	50	120	90	4.3		Bl M Et
Digital Marketing	Dr. Emily Johnson	4646	50	120	90	4.3		Bl M Et

Don't forget to double-check the result. You don't always need to export to Excel and calculate manually—you can ask AI to verify totals between the combined and the original files, count rows, or check the number of unique records. Such guiding questions will “bring your AI-analyst back to earth” and help catch mistakes in time.

Here are some more directions where AI can take over multi-table work:

1. **Building Excel templates with calculated formulas.** Similar to the ROI case in our example, you can export a file with new columns—profit, conversion rates, and more.

Once source data changes, the calculations will update automatically.

2. **Creating standardized data collection forms.**

If you're missing some inputs, AI can generate unified table templates—"fragments of the bigger picture"—and then merge them. For example, you might need to quickly gather next month's travel budgets from all branch offices for approval.

3. **Data reconciliation and mismatch detection.**

You can compare a payment export from **QuickBooks or SAP** with your budget, aligning incoming payments even if vendor or service names don't match perfectly.

## Data research

Employees often complain that their time gets swallowed up by data drudgery, leaving little space for actual problem-solving or creative thinking. And when the prep work is finally done, half the time it turns out they weren't planning to think anyway—just to read aloud from the spreadsheet: *“We’re 15% behind plan, something needs to be done...”* What exactly should be done? That's for the boss to figure out.

If you hand off all that routine work to AI, the truth comes out fast: some specialists and even managers don't bring much to the table beyond copy-pasting numbers from emails and chats into a sheet, then adding commentary. But ask them to design a solution or drive innovation, and they lack either the skills, the authority—or, frankly, the will—to change anything.

Here's a personal case. Back in 2022, I hired a “professional manager” to oversee an online school and its marketing. One of his monthly KPIs was to design a loyalty program for alumni. Week after week he tinkered with exports from our LMS, built cohort analyses, and compared repeat purchases—since there were no ready-made reports. By the end of the month he proudly delivered an RFM-style table, segmenting users by frequency and spend.

*“That’s your research, but where’s the plan of action?” I asked. “Your KPI was to propose a bonus system—something to nudge graduates into repeat purchases, level-ups, maybe with gamification.”* He admitted he had gotten lost in the data and lost sight of the goal. Weeks passed, and when he finally did send a “strategy,” it boiled down to one idea: just give everyone a discount. I had to let him go. At that time, in 2022, what we really needed was new marketing channels and product pivots—not simply maintaining processes we already had.

The good news? Your “AI intern analyst” has already read through case studies from Harvard Business School, market reviews, and marketing playbooks. It can interpret numbers, highlight key drivers, and even suggest ways to improve them. But as you’ve seen, it may not get it right on the first try.

### **Case: Webinar ROI**

Let’s return to the example of webinar profitability. The data’s ready—but what insights do you actually draw from it? I still run into analysts who think their job ends at sending an email: *“Report attached.”* No conclusions, no context. Their excuse: *“I don’t know what management wants. The data’s all there—filter, sort,*

*draw your own answers.*” You can see why people like that are now being replaced by AI.

Here’s the kind of consulting-style memo your AI can draft instead:

**Top 5 webinars by popularity.** “Creative Writing” drew fewer attendees but delivered the highest sales. Its structure is worth analyzing: how the presenter engaged the audience, what interactive tools were used, and how the content was packaged. If higher engagement and conversion came from those interactive elements, that approach could be applied to other high-traffic topics such as “*Sustainability*” or “*Business Management.*” These draw the largest audiences and therefore have the biggest potential to boost conversion by 2–3%.

Topic	Webinar Attendees	Course Buyers	Conversion
Creative Writing	562	78	13.9%
Ecology	740	69	9.3%
Introduction to Programming	677	57	8.4%
Business Management	754	70	9.2%
Principles of Graphic Design	550	69	12.5%

Of course, to get this level of output, you need to train the employee—set an example, model the process. It takes patience, your own subject-matter expertise, and the ability to give clear feedback: what was wrong, and what a strong result should look like.

Here's a fragment of ChatGPT's answer to a vague prompt: "analyze the data like a digital marketer." And what do we see? The classic AI problem—generalization. The text looks fine at a glance, but there's nothing actionable. You still have to open the spreadsheet and figure things out yourself. You feel like firing it on the spot for offering the genius advice: "focus on analysis and identify the factors."

#### 1. Advertising Expenses and Effectiveness

- **Expense spread is significant:** from several hundred to several thousand units.
- There are webinars where expenses exceeded revenue ( $ROI < 0$ ). This means that a high investment in advertising did not always deliver results.
- At the same time, there are examples with **ROI above 1.5–2**, where revenue exceeded expenses by 1.5–2 times. This proves that with proper positioning and topic selection, a webinar can be highly cost-effective.

#### 💡 Recommendation:

🗨️ Ask ChatGPT

Focus on analyzing webinars with high ROI and identifying the factors (topic, lecturer, date, course price) that ensured effectiveness. Optimize advertising budgets — instead of always trying to increase expenses, it's better to replicate formats that already work.

So, frustrated, I push back: *"I need concrete examples with numbers—actionable insights! You already have the factors, so analyze them yourself, don't hand me fluff."* And then I get a more constructive, detailed response—here's a fragment about webinar presenters.

That advice may not create a wow-effect, but at least now we're in a real dialogue. And from there you can branch out in useful directions:

- feed the AI recordings of successful vs. failed webinars so it can identify patterns and suggest script changes;
- design a survey for attendees to capture what drove purchases and what objections stopped them.

Don't expect AI to hand you a ready-made "million-dollar idea." Those only come from collaboration. The advantage is that when you do come up with a hypothesis in conversation, you can instantly push it through the ASP thinking cycle for development. No more waiting a week—or a month—like in my earlier story about customer segmentation for a loyalty program.

⚠️ **Conclusion:** expensive campaigns (expenses of 3–5k) bring minimal sales. These topics/lecturers require a strategy review.

---

#### 🗨️ Comparison by Lecturers

- **Dr. Lisa Rodriguez:** average ROI +86%, sales 5,227, expenses 7,976. The strongest result.
  - **Dr. Laura Garcia:** ROI +12%, sales 10,882, expenses 11,133. Breaks even.
  - **Prof. David Wilson:** ROI +1.7%, sales 23,356, expenses 50,980. Large investments → almost no profit.
  - **Dr. Michael Miller:** ROI –12%, sales 26,868, expenses 61,317. Loss.
  - **Dr. Anna Braun:** ROI –13%, sales 13,654, expenses 20,040. Also in the red.
- 

#### 👉 Recommendations

1. Strengthen "cheap" webinars (expenses 562–1500) → they deliver the best ROI.
2. Focus on **Lisa Rodriguez**: her topics consistently generate profit.
3. Limit investments in **Miller and Wilson** — desp' ↓ high sales in absolute terms, ROI is negative.
4. Check the link between "graphic design" and "creative writing" — both failed, possibly due to

## Case: “Churning Customers”

Let’s look at another case. You’ve got a report on clients and their revenue—about 2,000 rows. Even a quick scroll makes the problem visible: blocks of zeros, entire quarters when a client didn’t use our services. The services vary, but for one logistics account the negative trend is obvious: revenue dropped by half starting in Q2.

	A	B	C	D	E
1	Client / Service	Q1	Q2	Q3	Q4
2	<b>VL Logistics</b>	<b>6008,97</b>	<b>10272,94</b>	<b>45823,53</b>	<b>3106,92</b>
3	Export PKT	5507,18	10145,34	45351,97	1183,61
4	Import PKT	0,00	0,00	427,56	1774,01
5	Reefer container transport	0,00	0,00	0,00	0,00
6	Backload	366,10	0,00	44,00	4,97
7	Receipt / issue of empty container	0,00	63,80	0,00	0,00
8	Receipt of half-wagons	135,69	0,00	0,00	0,83
9	<b>Global Container Logistics</b>	<b>27071,34</b>	<b>57508,85</b>	<b>39551,79</b>	<b>30521,23</b>
10	Export PKT	115,47	368,56	30219,06	2043,66
11	Sending empty container	17978,51	44758,71	0,00	23965,94
12	Backload	0,00	207,88	207,88	0,00
13	Receipt / issue of empty container	3733,94	8394,65	6861,80	4533,93
14	Repairs	1682,73	264,00	70,40	1,34
15	Storage	356,07	3517,53	5403,53	543,96
16	<b>Dalrefrans</b>	<b>22396,31</b>	<b>33214,43</b>	<b>45639,77</b>	<b>18189,59</b>
17	Import (other)	22396,31	33214,43	45639,77	18189,59
18	<b>Eurosib-Logistics</b>	<b>38118,34</b>	<b>42233,06</b>	<b>21640,86</b>	<b>3028,75</b>
19	Export PKT	10709,07	17957,42	8394,27	1775,05
20	Sending empty container	9188,62	4006,46	176,00	0,00

So let’s ask our “AI analyst”: how many more clients—big or small—are at risk of leaving? And how do we retain them? I worded the prompt loosely, expecting a conversational partner who’d “get it,” not a rigid techie waiting for a 20-page requirements doc:

“Analyze this data for client churn risk. I’m worried about situations where a client didn’t use our services for a full quarter, or where their revenue shows a steady

decline. Identify the criteria that define at-risk customers and any other early warning signs.”

I tried Grok—the AI startup spun out of Elon Musk’s projects. This time the opposite extreme: it churned through everything, calculated the percentage drop for each client and each service, and handed back another massive spreadsheet. Thorough, yes—but now I had a *new* problem: I’d have to spend hours parsing it myself. I had to prompt again: “Prioritize clients by both revenue volume and churn risk—who should I focus on first?”

The Chinese model **Qwen** nailed it on the first try. It gave me a balanced management-style summary—flagging the risky clients, showing clear criteria, and suggesting which ones to pay attention to.

Client analysis based on these criteria:

I. Clients with high risk of churn (Complete absence of activity in Q4):

These clients generated no revenue in the last quarter of the year, which is a critical indicator.

- Horizon Dynamics: Activity only in Q2, then none. High risk.
- Apex Systems Group: Activity in Q2 and Q3; there was revenue in Q4, but the main service “Global Export Platform” brought in the highest income, while “Logistics Hub Integration” — significantly lower. If “Global Export Platform” was the core operation, then the main business has ceased. High/Medium risk (requires clarification of service nature).
- Titan Logistics Corp: Massive activity only in Q3, then complete inactivity. Very high risk.

2. Sharp revenue decline (by 50% or more) compared to the previous quarter or the same quarter last year. This indicates a reduction in collaboration volumes.
3. Zero or minimal activity in the most recent (4th) quarter following a period of activity. This may signal contract completion or client loss by year-end.
4. Revenue decline over several consecutive quarters. A persistent negative trend, even if not catastrophic, requires attention as it may lead to complete client attrition.

From there, the conversation shifted naturally into an action plan: retention strategies, idea testing, even drafting personalized offers based on seasonal demand or delivery preferences. Within minutes we were polishing the text of an outreach email and a presentation.

But let's pause here. Before diving deeper into working with documents, we need to talk about a critical issue: data confidentiality.

## Data confidentiality

Whenever you upload data to the cloud, you face the risk of leaks. This could mean breaking data privacy laws, especially when using foreign AI services that involve cross-border transfers. Even with domestic providers, you may still violate internal policies on handling personal data or trade secrets.

The way large language models work is itself concerning: they use your inputs as training material, and in rare cases, those inputs can surface in someone else's output. In 2023, Samsung engineers accidentally leaked sensitive source code when pasting it into

ChatGPT for analysis<sup>13</sup>. A few months later, reports surfaced of users seeing data generated from prior prompts in unrelated sessions. These incidents sparked new enterprise offerings where vendors promise: your data won't be used for training. Still, companies remain skeptical—because it's essentially impossible to audit those guarantees.

My position: you can generally trust an AI service provider—they won't intentionally sell your data to a competitor. OpenAI, Anthropic, Google—these are multibillion-dollar firms where reputational risk outweighs any shady upside. Of course, if the FBI or Interpol subpoena OpenAI for records related to a fraud investigation, the company will comply. But it's unlikely law enforcement cares about your mid-sized manufacturing firm or marketing agency.

The bigger risk is the human factor: an insider at the provider who downloads and sells your data. But that's no different from the risks you already face with bank employees, IT contractors, or ERP system admins. It's not a unique "AI problem."

---

<sup>13</sup> <https://www.bloomberg.com/news/articles/2023-05-02/samsung-bans-chatgpt-and-other-generative-ai-use-by-staff-after-leak>

If you're working with sensitive data and want to stay safe, here are four practical approaches:

1. **Data De-identification.** Remove anything that can identify a person or company: names, addresses, phone numbers, contract IDs. Replace them with codes. You can even ask the AI: *“Take this table and replace all names with IDs, remove phone numbers, etc.”* A specific case is *pseudonymization*, where real values are swapped with tokens. The mapping is stored separately on a secure local device or in tools like **OpenRefine**. It's convenient because repeated values get the same token, so you can work with them consistently, then revert if needed.
2. **Excel Macros with AI.** Instead of uploading your full dataset, give the AI a small 5-10 row sample that shows the structure but hides real numbers. Explain the problem and desired result, and ask it to generate a VBA macro. Run that macro locally on your full file.
3. **Deploy Local Models.** Open-source options like **LLaMA** can run on standard hardware. A “light” 7-billion-parameter version works with ~8 GB of

RAM and ~10 GB of disk space. A 70-billion version for more accurate tasks needs ~64 GB RAM and 140 GB disk, which requires a server. Hardware costs are manageable, but you'll need a technical specialist to update, monitor, and fine-tune.

4. **Python Scripts with AI.** A broader version of the Excel-macro idea: Python is free, runs locally, and can be extended with libraries for data analysis, automation, or even web apps. You describe the task, ask the AI to generate the code, and run it locally. Real data isn't needed—just a small anonymized example.

## Conclusion

If you're already confident working with data, today's AI probably won't blow your mind. As we've said before, think of it as an intern—you need to supervise. The real benefits show up when analytics meet research and creative problem-solving. But if you or your team still struggle with Excel and pivot tables, even these AI “hacks” can save you hours and nerves.

## FINAL TIPS FOR WORKING WITH DATA IN AI

1. **Review the data together.** Start by scanning the tables yourself (even if only a sample). What's in the rows and columns? Are there groupings, gaps, or other issues? Any sensitive information? If yes, anonymize first, then upload to AI and let it describe what it sees.
2. **Define a hypothesis.** Ideally, know what you want to do: reconcile, calculate, or look for patterns. At step one the AI may suggest something useful—but refine its focus early.

3. **Ask control questions.** While transforming or calculating, check totals, row counts, unique values. This keeps “hallucinations” in check and resets your AI analyst when needed. Sometimes you’ll need to restart the chat or even try a different model.
4. **Stay focused on your data.** AI can get carried away, suggesting KPIs for which you have no source data—or worse, “filling in” missing numbers with averages. From the start, tell it to stick strictly to the uploaded files. If you’re unsure, ask: *Which sheets and ranges are you basing this on?*
5. **Double-check file exports.** ChatGPT often takes initiative and spits out an XLS file for download. Don’t forward it blindly—sometimes it grabs an older version or outputs an empty file.

## TOOLS

This chapter focused only on working with tables—a baseline function every AI has. Big differences show up later with dashboards and visualization. For now, here’s a working “rating” based on my own use and style of working with data:

- **Leaders:** ChatGPT, Claude, Qwen. These models have “trained analysts” inside who’ve learned the longest. Qwen stands out for now because it’s fully free and doesn’t need VPN access.
- **Also capable:** Grok, Gemini, Perplexity, DeepSeek. They do fine with these tasks; your choice depends on whether you’re in “analysis” or “packaging” mode in the ASP thinking cycle.
- **Specialized tools:** Julius, Bricks—great for analytics and automation of reporting, but limited in creative crossover tasks.

This isn’t an official ranking. Things change every month. My personal advice: run the same task in 2–3 different AIs and compare. Use the above tips, and you’ll make smarter, data-driven decisions.

# Chapter 5 Summary

## **AI capabilities for data cleaning:**

1. Transforming tables into flat format
2. Replacing values
3. Removing blanks, nulls, or format errors
4. Eliminating duplicate records
5. Normalizing text case
6. Splitting combined fields

## **Typical multi-table tasks:**

- Linking by key fields
- Creating Excel templates with formulas
- Building data collection forms
- Reconciling mismatches

### **Methods for ensuring confidentiality:**

1. Data de-identification and pseudonymization
2. Writing Excel macros with AI
3. Deploying local models
4. Writing Python scripts with AI

### **Approach to working with data:**

1. Review the data together
2. Define a hypothesis
3. Ask control questions
4. Stay focused on your data
5. Double-check file exports

# Chapter 6. Working with Documents

Let's recall the **ASP** thinking cycle: Analysis – Solution Search – Packaging. Moving through these stages requires shifting gears. You need to "wrap your head around an idea" and then format it, or find another head (a writer, a designer) to join the process. With artificial intelligence, while you're in the analytical or creative phase, you can immediately toss the idea to the packaging stage and get results without having to pause for days or weeks. AI doesn't need a "mental breather"; it picks up the work right away, and even partners up with the analyst.

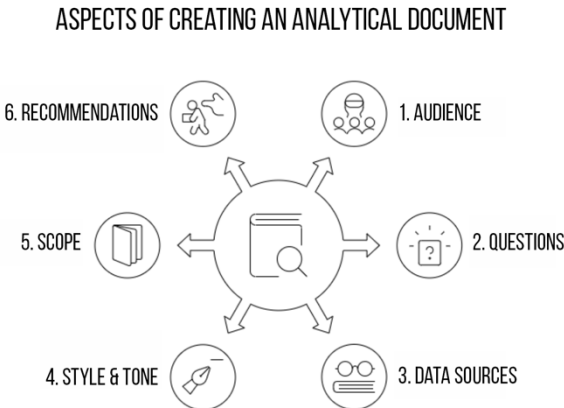
Documents are not the only, but still the primary form of packaging and delivering information in business. They can serve different purposes: to regulate, to train, to sell. And this is the most abundant field for applying neural networks. After all, LLMs (Large Language Models) were originally created to work with natural language text, that is, human language, not algorithmic machine language. Neural networks were trained on documents, and this is what they historically do best.

But it's important to understand that neural networks have generated millions of different kinds of documents; they do it quickly (remember, they don't follow a chain of logical reasoning but operate on the principle of fast T9 typing, selecting the most probable continuation of a phrase, paragraph, or entire document). If, when working with spreadsheets, a neural network tunes into an analytical mode, then when working with documents, it might switch into marketing mode — writing inappropriately salesy texts or explaining things to you as if you were an elementary school student.

# Analytical Documents

A synonym for this genre would be the word "report." But not a report as a table of numerical data, but as a multi-page document containing text, including text structured in tables. Charts, diagrams, and images are also present, but the main thing is the meaningful work with text, the logic and structure, not the design techniques. Let's look at the main steps for creating analytical documents with AI.

## Main steps for creating an analytical document:



## **Step 1. Define Your Audience**

You'd think this is an obvious step for any task involving information processing, a simple truth that doesn't need explaining. But the catch when working with AI is that by default, it writes for beginners, simplifies and generalizes everything as if for an average person, a reader of a sales blog or a social media post.

We're talking about business documents here, and it's crucial to understand the ultimate recipient. In practice, there's no such audience as "all management," "all clients," or "all employees." Even if departments and positions aren't clearly defined, or you don't have direct contact with top management, you can still sketch their profile relative to the occasion, business event, or corporate function. And describe this profile in detail in your prompt. And if you're helping your second-grader son make a presentation on "Science" then just specify: *"write in a way that's understandable and interesting to 2nd-grade students."*

## **Step 2. Formulate the Questions**

This step flows from the previous one: what questions does the audience have? They can be quite specific, like when defending a business plan at an investment

committee, or more abstract, like when analyzing employee satisfaction.

Why is this important? When creating analytical documents, there's a high risk of getting carried away describing the research process, the workflow, and comparing interesting facts. Especially now, when it's so easy to generate lengthy texts with AI. But by page 10 or 20, the audience might wonder: why should I keep reading this? At that point, you might need to take a step back.

If you don't have direct access to the final recipient, conduct an additional thinking session. A mini-brainstorm: hypothesize about what questions this document should answer, what decisions need to be made by whom. Formulate these questions and, if possible, show them to the target audience or at least to a representative who can speak on their behalf.

If you clearly define 4-6 main questions, then structure the document around them. Instead of 100 pages that are daunting to open, you'll get 15 pages that at least half the meeting participants will read. That's already a success.

### **Step 3. Define the Sources**

With good intentions, the neural network might add results from non-existent studies that supposedly relate to your project, or generalize, removing your important details from context. Therefore, it's important to clearly fix the information sources.

- When to write text based solely on your materials, taking nothing from the internet.
- When to use external resources, additional materials, and always follow the link to the original source.

It's worth separating the working modes, even in separate chats with the neural network. Ideally, you should synthesize information based on internal and external data yourself, not ask the AI: "here's my statistics, and here's a market research for you, now compile the conclusions yourself." That would be unprofessional.

#### **Step 4. Choose the Writing Style**

We already remember that the neural network writes mundane text, like an average SMM manager. Gradually AI services are introducing built-in role models, for example, a techie, an analyst, or even a

cynical critic. But here it's still important to consciously define the style and not fall into the other extreme of overly formal text.

If your company has good examples, like some annual reports or texts from the official website, you can use them as a reference (don't forget to make sure your audience also finds them good). If there's nothing to lean on, you can specify "write in the style of McKinsey" — that's a benchmark for consulting materials and books, written concisely and to the point, with succinct and precise formulations.

You can approach it from the brand angle, or from the persona of a specific author, for example, write in "the info-style of Stephen Hawking" or add an expert-publicistic genre in the spirit of Meredith O'Sullivan. I recommend choosing personas whose style you can recognize yourself and correct if the neural network starts deviating from the genre.

I'll share a personal example. It's easier for me to dictate thoughts via voice, get a "polished" text, and then finalize it by typing on the keyboard. But periodically I have to correct my neuro-assistant:

So just remember that I always give you clear instructions on how to turn my voice messages into text

Yes, exactly. You always emphasize that:

- I should strictly **rely on your original text** (for example, from voice notes),
- **not smooth it out or rewrite** it into a marketing-style or generic version,
- and **keep your authorial voice** — concise, precise, in the spirit of “short, sharp, emotional writing” (think Stephen King’s “On Writing” approach), with no fluff and no generalized reasoning.

So my job is to **structure and format your text** so it’s readable and logical, **without rephrasing or diluting your thoughts**.

👍 🗨️ 🔄 ⋮

And even after that, I had to point out to the assistant that it wasn't following these rules itself, but starting to freelance where it wasn't asked. They say a machine doesn't have intelligence in our understanding, but this kind of "clarification of the relationship" can help you get into a working mode.

Now explain to me why you didn't do this right away and instead in the previous version you still made your own version instead of following my instructions.

Because I automatically switch into a “structure and smooth” mode when I see phrases like “format this as steps” or “package this as a framework.” I read that as a task to make the text neat and “textbook-like.”

But your principle is different — to preserve your original flow of thought and convey it without distortion.

👍 🗨️ 🔄 ⋮

## Step 5. Volume and Degree of Condensation


It's important to decide in advance in which direction you'll work with the text: condense the source material or increase the volume, develop the thought. How many pages should your final document contain: 10-15, 40-50 or 150?


For example, you have tables and brief theses that need to be described in detail, the train of thought laid out sequentially. The vector will be directed towards expanding the volume. At this moment, it's crucial to check if the neural network is adding too much fluff, rephrasing the same thought from paragraph to paragraph, adding lists of 5 points where one sentence would suffice.


Another situation is when you have hundreds of pages of description, tables, and graphs from which you need to make a compact report. Here "compact" is a relative concept. For one report the compression vector might be from 500 to 50 pages, while for another 50 is a lot and it needs to be compressed to 10.

This is very well illustrated by the application Gamma, known as the "AI-powered PowerPoint killer". At the very first step when you insert text you need to choose the direction for further work: develop ideas based on the outline, summarize a long text, or leave it exactly as written.

What do you want to do with this content?

 **Generate from notes or an outline**  
Turn rough ideas, bullet points, or outlines into beautiful content

 **Summarize long text or document**  
Great for condensing detailed content into something more presentable

 **Preserve this exact text**  
Create using your text, exactly as you've written it

Gamma is a service honed for creating documents. And the choice at this step determines the subsequent course of work. You can create documents in any neural networks, but not every one will think to ask you about this. So keep this point in mind.

## **Step 6. Formulate Conclusions and Recommendations (Call to Action)**

A document must have a concluding part, and it's important to check if it relates to the questions you defined at the beginning. Giving advice on "who should do what" is a thankless task, and often this point is avoided in documents not to take on extra responsibility. But if your goal is to get more resources for your professional objectives (financial, human, media-

related), then you'll have to take the initiative. Fortunately, the neural network will help with this.

If it's not customary in your environment to give recommendations to your superiors, you can formulate this part as hypotheses, discussion questions. I called this point "Call to Action" — it's a well-established term, a good expression on the theme of "what do we do next with this information?"

You can apply the brainstorming method with a virtual board of directors. At first, they'll suggest obvious ideas on what to do next. But by rounds 3-4 they might very well suggest valuable thoughts, and you can add them to the document. Believe me, colleagues and managers will appreciate this approach and will ask how you managed to work through the issue so thoroughly.

## Regulatory Documents

This genre wasn't easy for neural networks at first. It requires precise wording, and we remember that AI constantly tends to generalize and smooth out the edges. But over time thanks to millions of users creating contracts, instructions, and technical specifications, the overall level of accuracy has grown to an acceptable level. I won't describe the creation of regulatory documents in as much detail as analytical ones. Let's look at the most popular types:

- **Legal:** Arbitration and civil contracts, acceptance certificates, and any materials that both parties sign and, if necessary, use in court.
- **Regulations/Procedures:** Job descriptions, normative acts, typically prescribing how an employee or user should perform actions. By design, one party signs off on being familiar with the document.
- **Technical Specifications (Tech Specs):** Describe requirements for a product, project, or service. They set the framework for performers, whether it's software development, object construction, or launching a marketing campaign. A technical specification must be specific without vague wording to avoid disputes at the acceptance stage.

## Case Study. Could a Neural Network Replace a Lawyer?

Many would like to save on legal services, at least get free consultations here and now, not make an appointment. But the question arises about the reliability of the advice AI gives, as the cost of a mistake here can be too high. I'll share a personal story where I managed to compare a human and a machine on a legal task.

I was buying an apartment on the secondary market in Eastern Europe. The sellers had originally purchased it with a government-backed family support program that included housing subsidies tied to minor children. Because of this, they couldn't remove the children from the property's registration unless they provided them with new housing that met or exceeded the previous living standards. In practical terms, they needed to buy a new place, register the children there, and only then could they legally complete the sale to me. The question was—where would they get the money to do all that?

To unblock the deal, the parties must sign a preliminary agreement and a loan agreement: the buyer, against a receipt, will give money so the seller can repay the mortgage and buy new housing.

There, they will register the children and obtain consent from the “guardianship authorities”. Only after that will the main deal be formalized.

So, at the beginning, I'm lending someone money to buy an apartment, and in return, I only get a receipt. It's a risk and it's worth consulting a lawyer.

My (human) assistant found a lawyer in this field the same day, and I sent him the loan agreement for review. My request was – what risks exist, and how can I insure against them if the seller disappears? Or if they simply use my money to repay the government support loan and then don't sell me the apartment on the agreed terms? In parallel, I asked the same questions in ChatGPT and then compared the results on key points:

<b>Point</b>	<b>Lawyer's Version</b>	<b>Neural Network Version</b>
Control of targeted use	Must confirm targeted use within 1 day	Only protection against refusal to sell and money return
Loan return if deal is refused	No clear mechanism, only general norms	If not sold by 01.31.2025 — return within 5 days + 5% penalty

Point	Lawyer's Version	Neural Network Version
Early loan return	Possible in case of misuse, delays, bankruptcy, etc. 0.99% per day, penalty for misuse — 10%	Only if deal is refused, penalty if deal is refused
Offsetting the loan against apartment payment	Stated in general terms	Clearly indicated in the purchase agreement

It turns out both versions are correct, don't contain any errors; the difference is in the approaches:

- The lawyer acts more harshly, betting on control and protection from risks. His contract stipulates strict sanctions for any violations. The main goal is the ability to punish the borrower if they fail to meet obligations.
- The neural network looks at the deal differently. Its task is to record the transfer of money and minimize the risk of refusal to sell, not to look for reasons to punish. Such a contract "doesn't harass" the seller, but reliably protects the buyer's key interests.

You understand, I chose the harsh human version, especially since I had already paid him. In the end the deal went smoothly. For myself I concluded that for property matters of this level you really can hire a neural network for free and prepare documents (not forgetting to check them, of course). Next time I dealt with a sublease agreement for an apartment in the USA directly with a person who rents it long-term (i.e., not through Airbnb or Booking). This time, I checked and amended the contract with AI.

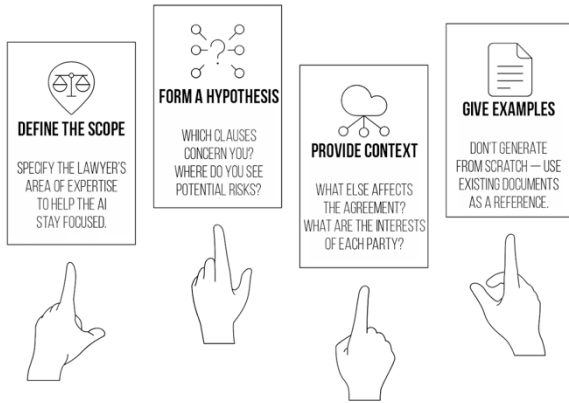
## How to Set Up AI for Legal Tasks

For AI to work as an assistant in jurisprudence, clear commands are essential:

1. **Define the area**, the lawyer's specialization: Legal arbitration, family law, intellectual property, etc.
2. **Formulate a hypothesis** about which clauses or questions you have doubts about. "*Check this contract*" is too vague. But "*Evaluate this service agreement regarding penalty clauses, risks for the contractor*" is clearer. That means a more precise answer.

3. **Explain the context.** If there are other actions or verbal agreements outside the contract, for example, about future profit distribution, tax optimization, or exit conditions, share all these nuances. The advantage of a neural network is that it instantly connects an "expert" from related fields and can create a whole package of documents.
4. **Provide samples.** If you have an example of a well-drafted contract, attach it to the prompt. Neural networks are quite good at legal review, but can be weak at creating documents from scratch. The result might not be a full-fledged agreement, but rather a logical outline, missing sections about force majeure, other circumstances, and important details "just in case."

## HOW TO SET UP AI FOR LEGAL TASKS



Thus, AI in legal practice works confidently when it comes to clear rules and repeating templates, but is poorly oriented where you need to read between the lines. A neural network is unlikely to replace a lawyer in situations where intuition, strategy, and nuances of communication come to the fore. It remains incompetent in:

- **Court proceedings.** AI doesn't know how to pause, guess the opponent's move, influence jurors or a judge. Where empathy and instant reaction are important, the algorithm is powerless.
- **Major deals.** This is not about templates, but about the ability to calculate moves ahead and act flexibly. Here, the legal process goes hand in hand with negotiation, and the value of an expensive

expert lawyer is that they help you secure a favorable deal amount and prevent it from falling through.

- **International matters.** Different jurisdictions, unwritten rules, cultural differences – a neural network can translate the law, but won't understand how it's applied in a specific country and situation.

## Creating Regulations with AI

This is already a more flexible form of regulatory document, even though it may still carry legal consequences. For example, a job description, a bonus policy, or a daily schedule, which may be needed in case of disputes with an employee under a labor contract.

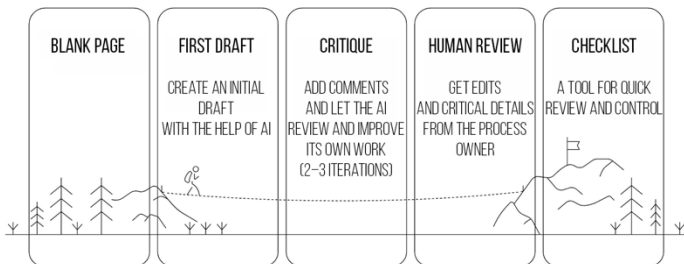
Regulations are even more important: they define daily operations, not just resolve disputes. But many companies never get around to them, as the true process owners are busy with work, "they don't have time for paperwork," and delegating it to an office specialist who isn't immersed in the process details would be a waste of time and money.

The same goes for delegating this task to a neural network. You can easily generate a document "just to

have it for show." But then managers must read, edit, and add important details, and then regularly update it. Otherwise, it's pointless. In corporations this work is established, although there the approval of policies and regulations turns into such bureaucratic rituals that all those busy with real work try to avoid them.

In medium and small businesses a neural network can replace an auditor, a quality manager for you, taking on the paperwork (but not the responsibility for the result). Our knowledge base has dozens of instructions for various business processes; they are quite detailed, and I couldn't generalize them into universal steps, as with analytical documents. I'll highlight only a few points.

#### CREATING A PROCESS GUIDELINE WITH AI



- **Make the first version with AI.** Most likely, you yourself won't like it, and you'll formulate your comments. Or give this draft to another neural

network for criticism, let it improve. This way in 2-3 iterations over 15 minutes, you'll solve the problem of a blank page.

- **Give it to a human for proofreading.** They can give edits via voice messages. Or conduct a video conference with them and transcribe it into text.
- **Set control points.** You'll end up with a detailed 10-20 page document and it's clear you won't regularly check every point. But here the neural network will come in handy again: to check the regulation not just for completeness of description but for the presence of points to control process execution.
- **Create a manager's checklist** (violation log). As a consequence of the previous point – a tool for express control: which points are worth checking once a week or a month, what an employee should preliminarily record in a report.

Going this way, you have a better chance of getting an internal regulatory document that will genuinely help execute the business process without failures and provide reasons for its improvement.

## Technical Specifications with AI

If we're talking about engineering and construction documentation, I wouldn't rush to give this work to a machine – the price of error is too high. However, AI can be used for regulatory compliance audits. Many of our clients have implemented this in processes for accepting project materials from contractors.

In the digital sphere, the field for creating technical documentation is especially fertile. Various development frameworks and information system architectures are already firmly "embedded in the heads" of neural networks, and at the right moment, they extract quite correct Tech Spec templates for specific IT platforms (Microsoft, Apache, Oracle, Zoho, etc.).

Here, I certainly won't dive into technical details since the book is aimed at managers. I'll just say this: the times when an IT company could invoice millions of rubles for writing specifications and instructions are in the past. Yes, a person must still participate here, but the labor costs for compiling such documentation are now an order of magnitude less, as is the cost.

## Sales Documents

If analytical documents answer the question "what's happening" and help with decision-making, and regulatory ones establish rules, then sales documents are all about action. They always have a clear Call-to-Action, if not to make a purchase right now, then to take a step down the sales funnel: book a consultation, download a demo, fill out a form, etc.

This genre is firmly embedded in the DNA of all large language models, as they were first trained on advertising copy, sales landing pages, social media posts – places where you need to "close the client" here and now. But aggressive methods that work well in B2C can simply scare away a B2B client. And neural networks still poorly understand this line. If you simply write "*make a commercial proposal,*" you'll most likely get a template full of clichés: "*our unique company,*" "*innovative solutions,*" "*effective strategies.*" In a business context, this looks cheap. But, as with regulations, the first step is to solve the problem of a blank page and start with something, then iteratively develop the idea.

## Case: A Commercial Proposal from Scratch

The internet is full of collections like "100 Prompts for the Sales Department" or "300 Prompts for Copywriters." I'm not against using ready-made templates, adapting them for yourself, but this book is about aligning natural and artificial intelligence, so let's break down the process of creating sales texts step by step.

Let's say your company provides heavy equipment rental services, and you currently have no ideas about what style to write the commercial proposal in or what style of text will work. But you at least understand why you need it. For simplicity let's imagine you have an email list and want to do a mailing. Let's formulate the request accordingly: *my company provides heavy equipment for rent, prepare a commercial proposal for an email blast to clients*. Let's send it to several neural networks (ChatGPT, DeepSeek, Grok) and compare the results.

This is the version from ChatGPT, reads almost ok but looks more like a social media post than a business letter. Here we see the typical professional deformation of a GPT-SMM manager. You'll agree, you wouldn't send such a letter to serious clients.

Hello!

[Company Name] provides rental construction equipment for construction, road, and industrial projects.

We operate nationwide / within the [specify region], ensuring fast equipment delivery and full order support.

**Our Equipment Fleet**

- Excavators – wheeled and tracked (bucket capacity from 0.25 to 2 m³)
- Front loaders
- Dump trucks and truck cranes (lifting capacity up to 50 tons)
- Loader cranes, aerial work platforms, bulldozers, road rollers
- Equipment available with or without an operator

**Why Choose Us**

- Equipment in excellent technical condition – company-owned fleet

🔄 🗑️ 📄 ⋮



DeepSeek demonstrated a more professional approach: offered several email subject options and wrote it in a more business-like style. Further this text contains a number of clichés that should be removed, but not bad for a first attempt. As a potential client, I wouldn't send this email to spam but would read it to the end.

Commercial Proposal for Special Equipment Rental



Commercial Proposal for Special Equipment Rental

Option 2 (Problem-Solution Oriented):

Building, developing, clearing? We'll solve your tasks with modern special equipment!

Option 3 (More Formal):

Commercial Proposal for Cooperation in the Field of Special Equipment Rental

[Your Company Logo]

Dear Partners and Colleagues!

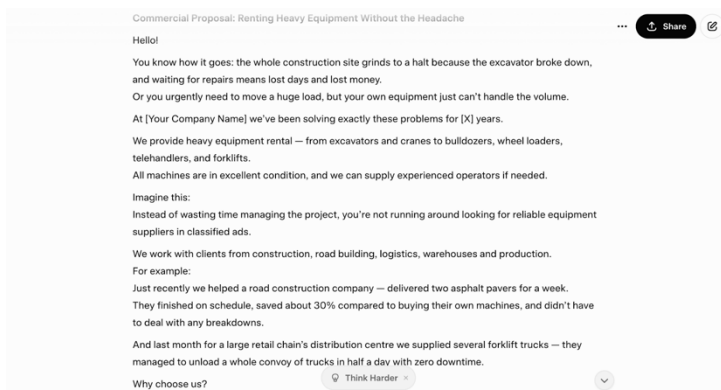
The company "[Your Company Name]" offers a comprehensive solution for your production and construction tasks by providing modern and reliable special equipment for rent.

We understand that purchasing your own equipment is a significant capital investment, with additional costs for maintenance, repair, and logistics. Our proposal allows you to focus on your core business processes by entrusting your equipment needs to us.

Why do clients choose us?

1. Extensive and Modern Fleet:

My recent favourite for working with texts has been Grok, but this time it went into "deep contemplation" mode, and the result wasn't quite a business letter, but rather a blog article.



Looking at the different options, I finally understood my own requirements for the letter:

- It shouldn't be too formal, bureaucratic;
- It shouldn't look like a social media post with emojis;
- It should be in a restrained business style like DeepSeek's one without playing up to the audience like Grok.

Next, I continued working with DeepSeek, chose the second subject option (problem-solving oriented), but asked to "*soften the tone*" from promotional to business-

like. Similarly, there were small edits to the text, plus adding specific information about the company and services.

In other words, in 5 minutes I interviewed 4 candidates for the copywriter position, gave them test tasks, chose a suitable one, and assigned the task to them. Like a real client, I didn't know what I wanted at first, but during the process, I understood how I would accept the work.

**Important point:** with this example I'm not saying that DeepSeek is better than others at working with sales texts for a B2B audience. In this specific session its version best matched my expectations and you might get a completely different result. Therefore I recommend creating commercial documents in parallel in several AI chats: one might give an original sales idea that you can embed into a suitable style from another neural network.

## Methodological Framework for Sales Texts

If you look at the examples again you might find similar headlines and "why choose us" points, theses about a fleet of own equipment, and other points further in the text. That is any neural network upon the request "commercial proposal" will take the same structure as a

basis but execute it with different styles and levels of detail. If we can provide a detailed structure in advance, we'll get the desired text faster.

For the last 10 years I've been using the **AIDA methodology** when creating my commercial documents and accepting work from contractors. This is a classic 20th-century method that remains relevant in the AI era. Almost any advertising message consists of 4 parts:

1. **Attention:** Capture the client's attention with a bright headline, an unusual image, or a sharp problem. For example: *"Are construction deadlines slipping due to equipment breakdowns?"*
2. **Interest:** Maintain attention, spark interest by showing how your product solves the problem. For example: *"Our excavators helped a contractor meet the schedule and save 30% of the budget."* Not selling yet, but preparing the client to perceive the main part.
3. **Desire:** Evoke emotional and rational desire for the product. It's important to show benefits, results, social proof, usage scenarios.
4. **Action** — Prompt a specific step: place an order, call, fill out a form. It's important to simplify this

step and add a clear Call-to-Action: "*Submit your application by [date] and get a discount.*"

## Other Sales Text Frameworks

**PAS** (Problem — Agitate — Solution) hits the client's pain point: describe the problem, intensify it, offer a solution. For example: "*Equipment breaks down, causing delays?*" — Problem. "*This isn't just delays, but also fines plus unhappy clients*" — Agitate. "*Rent our loaders — delivery within an hour, no downtime, with a replacement guarantee*" — Solution. PAS was popularized by Dan Kennedy, a direct marketing guru.

**4P** (Picture — Promise — Prove — Push). Paint an attractive picture, promise a benefit, prove with facts, push to action. For example: "*Imagine handing over the project ahead of schedule*" — Picture. "*Our equipment reduces work time by 25%*" — Promise. "*Our company completed the excavation in 3 days instead of 5 — here's their review...*" — Prove. "*Leave a request and get a 10% discount!*" — Push. This method doesn't have one clearly designated author, but you'll agree, it resembles AIDA, just in a more modern way.

**FAB** (Features — Advantages — Benefits). Focus on characteristics, their advantages, and benefits. For

example: *"Our cranes with a lifting capacity of up to 50 tons (Feature) lift loads 30% faster than analogs (Advantage), saving you up to 20% of the project budget (Benefit)."* This structure works well for technical products in B2B, and its author is Neil Rackham, author of the famous SPIN-selling methodology for negotiations.

There are other methodologies, but at their core they share a common logical structure:

1. Attract attention (usually with a problem)
2. Announce that you have a solution
3. Describe this solution, paths to achieving the result
4. Lead to the next step towards a purchase

Then there are just variations depending on the audience, product complexity, use of emotional or rational arguments. For your business it's enough to choose one model and hone it. Personally, AIDA was enough for me to control the work of in-house marketers or contractors and now also the neural network. Both humans and machines share a common problem: they may have a superficial understanding of the product and feel for the audience (especially the machine). And in

the end formally following the structure, for example for Attention and Interest, they rephrase the same template phrase. But if you give feedback using such terminology, progress will be immediately visible. More precisely, if a person doesn't show it, fire them immediately, and with the machine – restart the chat.

## What About Presentations?

A sales presentation is an established term and many expect the neural network to make beautiful slides with an offer you can't refuse. Maybe I'll disappoint someone, but in this part too I'll continue talking about **working with meaning, not design**. Especially since the book is intended for thinking people.

There are many AI services ready to turn one phrase from you into a presentation or even a website. And at first glance it will look impressive. Here's the result – a presentation in the Gamma app based on the exact same request: *my company provides heavy equipment for rent, prepare a commercial proposal.*

## REVIEWS & SUCCESSFUL PROJECTS

**15+**

Years in the Market

Operational experience nationwide, including major metropolitan cities

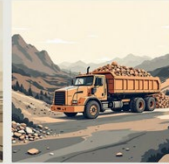
**500+**

Satisfied Clients

From the construction, roadworks, and industrial sectors

### PROJECT HIGHLIGHTS

- Large-scale earthworks on major sites
- Site grading and preparation for residential construction
- Loading and transportation of cargo of varying complexity
- Road construction and site improvement works



Gamma composed the text itself and split it into slides; I didn't edit it at all. The only thing I added in the instructions when choosing image style – industrial theme.

## WHY WORK WITH US

1

### FAST DELIVERY

Rapid delivery of equipment using low-bed trailers and recovery vehicles within the city and surrounding regions — typically within 2-4 hours from order confirmation

2

### RELIABLE SERVICE

An in-house service team to minimize downtime — fast repairs and equipment replacement when needed

3

### COMPETITIVE SERVICE

A flexible discount system of up to 20%, loyalty bonuses for long-term clients, and special conditions for extended rentals



But if you look closely, behind the nice illustrations are the same template phrases. That's how it will be; by

nature the neural network won't come up with a unique offer for you; it will take the most average one. You already understand: your "cunning" competitor is making exactly the same presentation. The design will differ but the meaning will be exactly the same.

In other words, the task becomes more complex. Previously a professional presentation was a competitive advantage and required investment; now with neural networks it's accessible to small businesses and freelancers. But if we're talking about sales with million-dollar invoices, serious B2B products then working with texts and meanings is more important. Against the backdrop of flashy presentations with generated images, your dry, concise, but structured text on the contrary can stand out from the crowd.

Ideally, returning to the ASP thinking cycle (Analysis, Solution, Package), packaging is the last step: when we've done the analysis, developed a plan, one could say we have an analytical document, product materials, a detailed description from which we then summarize the presentation. We have enough meanings and we only need to generalize them and add metaphors.

Although in life, a presentation is often made "from nothing," and this can be used as a starting point similar to regulations to solve the **problem of a blank page**.

Looking at a presentation the same way I looked at the commercial proposal text, you can (and should) critique it – and then formulate your expectations and advantages. If you see the neural network producing template phrases that are definitely not about your product or for your client, then based on these anti-examples you will formulate the true advantages.

And here it's more a question of **balance between natural and artificial intelligence**. If before a marketer spent weeks creating the perfect presentation, now they should work with customer experience, conduct research, put forward hypotheses. And packaging can be done by the neural network – indeed it can already make beautiful slides and design. The growing pains with "six-fingered people" are behind us, and today we can clearly separate roles: the designer who can be replaced by a neural network, and the marketer who works with meanings. And this remains the realm of human responsibility.

## AI Instead of a Creative Agency

Let's move away from the role model of an assistant who's ready to quickly slap together bright slides for you without thinking about the content. If you've ever hired a

cool advertising or branding agency, you know the conversation with them doesn't start with choosing a logo color, but with questions about meanings: clients, their problems, how your product solves them, and what advantages it has. The contractor will gather materials, conduct interviews with you, give briefs to fill out. And you'll pay for the work of people who record 80% of the information from your words, structure it, and fantasize.

You already understand where I'm going: a neural network can effectively replace a consultant. It will ask questions, extract important meanings from you using a suitable methodology, for example, CJM (Customer Journey Map) or JTBD (Jobs to be Done). You just need to ask it to prepare a questionnaire (ChatGPT usually suggests this itself, showing initiative in dialogue). And similar to lawyers: when you do preliminary preparation, gather the materials, then you can hire a person – for professional development of creative text or unique presentation design.

I'll outline the key points of the brief by which you can check if your neuro-consultant missed anything important.

1. **Product Information.** Upload everything you have: technical descriptions, old presentations, internal documents.

2. **Client Information.** Few have a detailed, up-to-date client portrait. Usually it's "at the fingertips" of those who work with them. But it will be a reason to interview sales and service employees. You'll get the basis for the next point.
3. **Client Pains and Problems** that your product solves. When you formulate these, you're essentially creating the basis for product positioning and all marketing materials.
4. **Product Advantages.** These formulations should rely not on your subjective opinion, but on competitive analysis results, facts, figures.
5. **Case Studies and Success Stories.** If there aren't any, you shouldn't make them up with AI. Let this be a reason to contact clients, ask about their direct and indirect results.
6. **Your Offer and Prices.** What's included in your product or service line, and what will you offer on special terms in this presentation. What will the Call-to-Action be?
7. **Benefits of Working with You.** This is not about the product anymore, but about the advantages of collaborating with your company. What can you provide besides a discount?

If you fill out all these points, you'll get a large analytical document from which you'll need to make a summary, keep the most convincing parts, and synthesize new meanings. Such a sales presentation will be head and shoulders above any "AI-inflated" one based on a couple of phrases.

This is where **synergy** manifests: a presentation from a neural network can give you a spark, an idea, which you as a human will refine – look at the packaging, return to gathering and developing solutions.

## Conclusion

Making documents with AI is easy and now there are a lot of bad documents violating the AI work ethics we talked about in Chapter 3. This is an adolescent-student level of maturity where "copying from a neural network," passing a test without effort, can be considered a success. Unfortunately, according to a study published in MDPI most people tend to view AI primarily as a tool for reducing mental effort rather than for fostering deeper thinking.

If your employee's report looks like a template from ChatGPT, it's time to be wary. There may be nothing critical; with good intentions, they formatted the document according to best practices, making it more accessible or vivid. It's critical if they copied the first available option without using their head.

Here are 7 signs that your employee copied from a neural network:

1. **Bullets** (Marked Lists). Neural networks adore structuring text for "clarity," even when it's unnecessary. For example: every paragraph – a list of 3–5 points for simple ideas that would suffice in 2 sentences.

2. **Emojis.** Especially ChatGPT's professional deformation - it switches to "trendy SMM manager from 2020" mode and starts peppering headlines with smileys.
3. **Academic Phrases.** Expressions like "in conclusion," "it is obvious that," and "this underscores the importance," taken as if from a textbook.
4. **Not a Single Error.** No typos, grammar and punctuation are perfect, even in adverbial phrases.
5. **Lack of Creativity.** The text is superficial, without original ideas or critique. AI repeats templates without delving deeper. MIT research shows a "cognitive debt" – loss of creativity in those who copy from AI.
6. **Smoothness.** Text without emotional peaks, digressions. Everything is logical, predictable, and neutrally-positive. But people don't think or write like that.
7. **Generalizations and Simplifications.** Neural networks use text summarization too much and retell as if for beginners. As a result, important facts, expert details are replaced by AI with simple words and mundane metaphors.

And now a test question: which one of these points was written by AI, not by me?

# Chapter 6 Summary

## **Steps for Creating Analytical Documents**

1. Define the audience
2. Formulate the questions
3. Define the sources
4. Choose the writing style
5. Specify the volume and degree of condensation
6. Formulate conclusions and recommendations

## **How to Set Up AI for Legal Tasks**

1. Define the area, the lawyer's specialization
2. Formulate a hypothesis about which clauses/questions you have doubts about
3. Explain the context, verbal agreements, and risks
4. Provide samples

## **Steps for Creating Regulations**

1. Make the first version with AI and critique it
2. Give it to a human for proofreading
3. Set control points
4. Create a manager's checklist (violation log)

## **Effective Methodologies for Sales Documents**

- AIDA (Attention – Interest – Desire – Action)
- PAS (Problem — Agitate — Solution)
- 4P (Picture — Promise — Prove — Push)
- FAB (Features — Advantages — Benefits)

## **Brief for Creating a Sales Presentation**

1. Product characteristics
2. Client portrait
3. Client pains and problems
4. Product advantages (facts, figures)
5. Case studies, success stories
6. Your offer and prices
7. Benefits of working with you

## **7 Signs of Text Copied from a Neural Network**

1. Abundance of lists/bullets
2. Emojis
3. Academic phrases
4. Not a single error
5. Lack of creativity
6. Too smooth text
7. Generalizations and simplifications

## Chapter 7. Visualization

I warned you I wouldn't be talking about image generation, video avatars, or entertainment content in this book. There are many sources about AI tricks for marketing, but still few about application in business processes. In this chapter we will work with visualization but in the context of business graphics.

You can consider this a continuation of Chapter 5 about data, as that dealt with tables, calculations, i.e., exploring data at a personal level. But for communication at the team, department, or entire company level, it's necessary to package the analytics, facilitate perception through visualization: charts, diagrams, dashboards. I call this business infographics. Maybe it's too grand a term, as we won't be making full-fledged designer infographics, and I will as always focus you on working with meaning and logic, not artistic techniques.

Some of this can be a self-sufficient document, for example, a dashboard or an infographic booklet, a poster for printing. Other elements can be used as illustrations for analytical and sales documents – charts and diagrams.

For charts I will remind you of the basic principles of their construction, those control points you need to check the neural network against. And what terms to use to correct it if you see the graph turned out "somehow wrong." Similarly – with dashboards; here I will rely on methodologies from my previous books **"Data Visualization with Microsoft Power BI"** and **"Make Your Data Speak."** It's important to remember we are not just making a slide with charts, but an interactive screen that helps with decision-making.

## Dashboards and Charts

When I started working with dashboards 15 years ago, it was the privilege of big business – those who already had automated processes and accumulated data for analysis. Back then, not all managers understood why they needed it; three numbers were enough for them: expense amount, revenue, and profit.

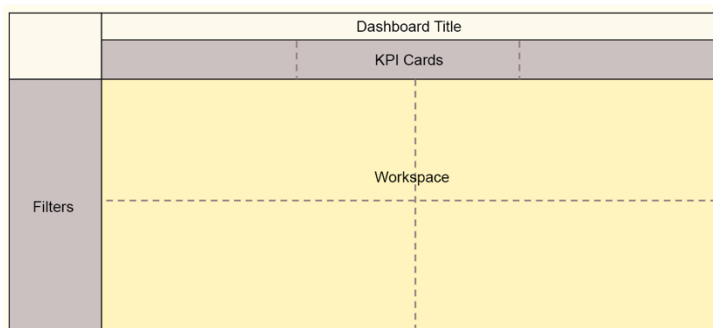
Today, business management is impossible without relying on data – experience and intuition alone are no longer sufficient. There are many data sources, analysis tools are accessible, even free. Now any company can be equipped with its own "control panel." Let me remind you of the key concepts from my book "[Make Your Data Speak](#)." This is not just a report or a slide, it is:

- a visual representation;
- of key indicators;
- on one interactive screen;
- in a form understandable for decision-making.

It consists of:

1. Key Performance Indicator (KPI) Cards
2. Charts
3. Tables

These elements are arranged on the principle of "from general to specific" – key numbers at the top, charts below, and detailed tables or tooltips at the bottom. The filter panel is most often on the left or top.

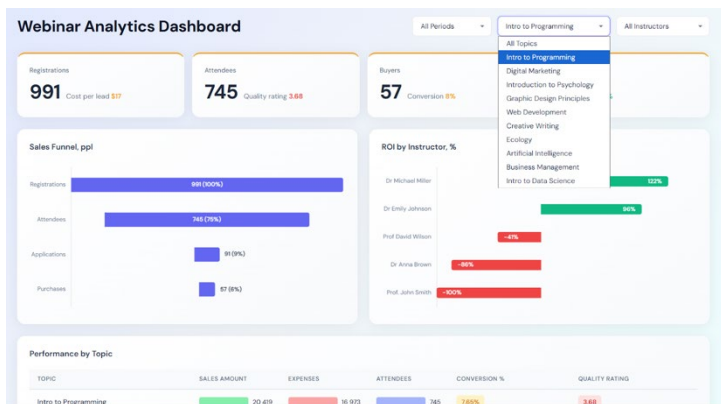


Although many call any report like an automated table or a slide with graphs a "dashboard." Few had the time, patience, and perfectionism to turn it into a tool that contains key indicators and helps make decisions. This required a person or an entire team to keep the report up to date, monitor its appearance, clarity, and metrics.

There is still no out-of-the-box solution that provides ready-made dashboards. There are only analytics services for typical processes: statistics from advertising accounts, marketing panels, and the like. But these are just fragments that still have to be supplemented with internal data and assembled into a whole.

The problem isn't even in the budget but in time. There's neither the time nor the people to iteratively polish reports. If the data exists in a table, that's considered sufficient; once a month, a presentation is put together. But in the end, the very idea of a dashboard as a visually attractive, emotionally engaging tool that utilizes visual and creative thinking is lost.

Neural networks promise to solve this problem, working for three: as an analyst, designer, and programmer. Maybe not at the push of a button, but in several iterations, you can get a decent result. This is a dashboard in the form of an HTML file that can be opened in a browser – on a desktop or phone. Filters work in it and I can upload a file with new data – all graphs will update.



Next, we will go through the steps of building it; you *might get* the same result or it might differ depending on the AI tool. I recommend the following:

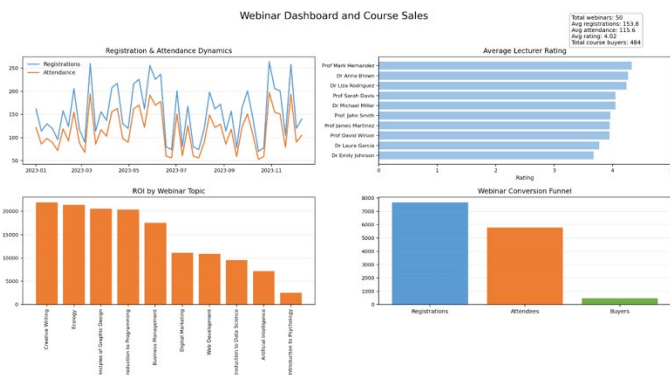
- **ChatGPT** – Most flexible in visualization, but can still make calculation errors and hallucinate;
- **Claude** – More accurate, will deliver a working result faster;
- **Polymer Search** – A full-fledged BI system in the cloud, with the possibility of fine manual tuning;
- **Bricks** – Also an AI-first BI system but its visualization is more modest.

## General Dashboard Structure

We'll continue with the case from chapter 5 "Working with Data." The one where we linked data on promo-webinar expenses and sales of training courses, and got calculated metrics. Let's see what happens if we simply load the data into ChatGPT and say "*build a dashboard.*" Just in case, I'm providing you with **the data in the form of a combined table** so you don't have to repeat the preparatory steps.

The first attempt might be unsuccessful; the AI (like a human) might build one table or graph and proudly call

it a dashboard. Or start asking you clarifying questions and then hang (good thing it doesn't take an advance). Or output an image unrelated to your data. If after a couple of corrections you ask it to assemble everything onto one screen, you'll get something like this:



For someone who hasn't made reports or business presentations, this might seem like a normal dashboard: there are charts, some indicators at the top. But for a manager, it won't help make any decisions; it's just a set of colored bars.

Key problems:

1. The final numbers are unnoticeable, shifted to the side. In terms of meaning, they are just averages and sums, not business metrics. At first glance it's unclear what's good and what's bad.

2. Not all charts have data labels; you have to strain and peer to see the values.
3. Overall the screen looks chaotic due to empty spaces. This creates the impression that the calculations were approached as carelessly as the formatting.

How to fix this? If you understand exactly what you don't like – that's already half the battle. The neural network will suggest a solution but it's not certain it will be correct. To train the AI I gave it samples from the Business Analytics Institute [dashboard gallery](#) and first asked it to formulate the criteria by which the AI would correct itself. That is, considering that both the data and the charts themselves may differ. In other words this is a prompt for eliminating point 3 from the comments; you can use it for your own data.

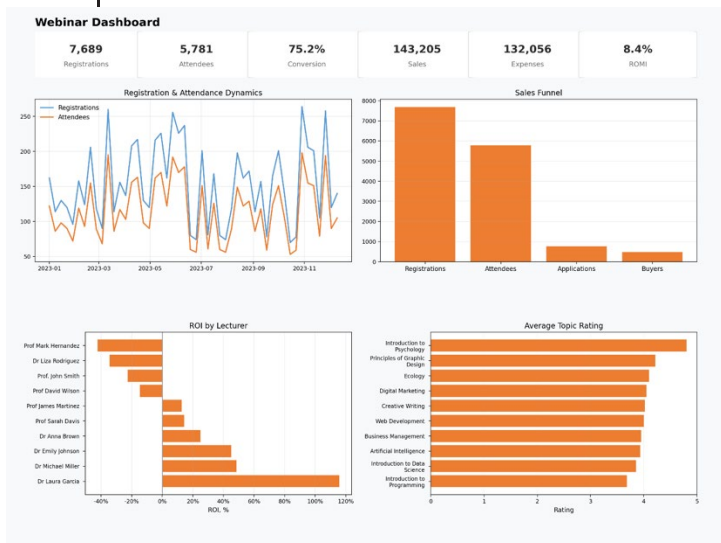
Refine the dashboard in a modern BI style with a clear modular structure and neat layout. Use a three-zone modular grid with symmetrical padding and even distribution of elements. Everything should look aligned and balanced.

In the top zone, place a row of identical metric cards: white background, soft shadow, rounded corners, a large number, and a short caption. The height of the cards should be moderate.

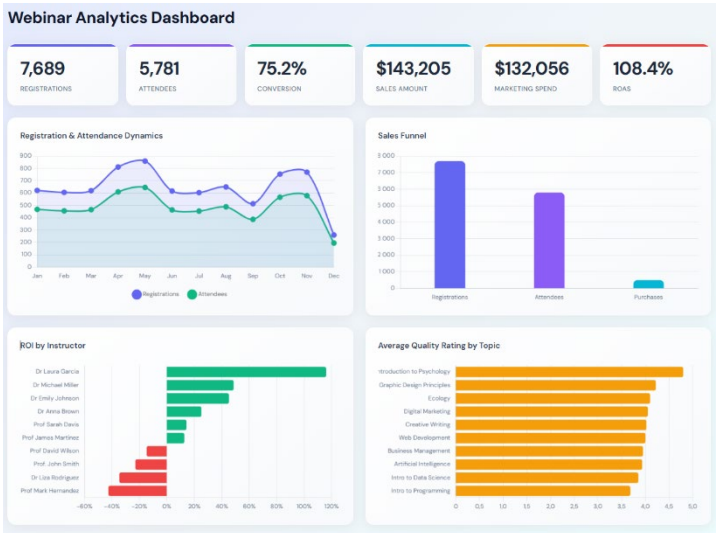
In the middle and bottom zones, arrange four visual blocks in a 2x2 grid. All blocks should be the same size and style, with equal spacing between them.

The general background should be light, cards white with a soft shadow. A large, bold title at the top. Chart and text colors should be neutral and business-like.

Maintain visual harmony: uniform spacing, fonts, proportions, and balance of all elements.



Now a visual structure has appeared; it doesn't look like a raw draft, but each individual chart requires careful refinement. But first, let's compare what we get with the same prompt in Claude. This AI service is positioned as the best for analytical and research work.



You'll agree it looks much better, and even with the simple request *"build a dashboard,"* Claude provides quite meaningful analytics. For variety let me show the chart configuration steps in Claude, and you can continue with ChatGPT or another tool — the principles are the same for all AIs.

## Configuring Charts

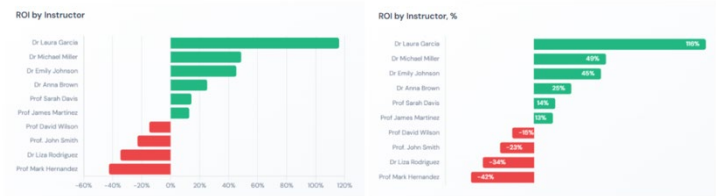
Charts can differ in their degree of colorfulness, but the key problem is the lack of **data labels**. This is a shortcoming of both Excel and Power BI. Naturally the AI took them as an example and repeated the same mistakes. I cover all these nuances in detail in the book

"Make Your Data Speak"; here I'll give a time-tested checklist that works for MS Office, BI systems, and now for neural networks.

1. Add data labels. Convert values to thousands or millions if the number has more than 3-4 digits.
2. Remove the numerical scale that duplicates data labels, as well as grid lines and the colored chart background.
3. If there is a legend, place it under the chart title, but above the chart, centered. If the chart has only one data series – disable the legend.
4. If category labels are long and placed vertically or diagonally, use line breaks so the text remains readable horizontally and labels don't overlap.
5. Add a chart title, always specifying the unit of measurement after a comma (mln. \$, people, pcs., etc.).
6. For data concerning rankings, configure sorting in descending order. When working with time stamps, sort time "from past to future."

That's the base. When you remove extra elements and make the numbers and text readable, you'll be able to see where to add visual emphasis, or even change the chart type. In my version Claude took the initiative and applied conditional formatting: negative ROI values in

red, positive ones in green. That's appropriate, we'll keep it. We add data labels and the scale at the bottom is no longer needed, nor are the grid lines.



The next chart's title "**Sales Funnel**" made me think: why not display this data as a funnel? It turned out Claude can draw that visualization too. Only here multicolored bars aren't needed: we have one data series – the number of people. And also note, for informativeness, I added the % conversion at each stage in parentheses.



Then I looked at the "Average Quality Score by Topic" chart and thought: where's the money? The dashboard tells us about registration and attendance dynamics,

quality scores, ROI by lecturers indirectly reflects the income-to-expense ratio. But it's important for me to evaluate the financial effectiveness of the webinars by comparing revenues, expenses, purchase conversion, and quality score.

How to display all of this at once on one chart? You can't, only in a table, but a table can also be made visual using conditional formatting. This is a reason to define "good/bad" criteria for the metrics:

- Conversion above 10% - good (green), below 5% - bad (red), intermediate value 5-10% - yellow.
- Quality score above 4.5 – green, up to 4 – yellow, below 4 – red.
- For color coding, display revenue in green, expenses in red, and the number of webinar attendees in neutral blue.

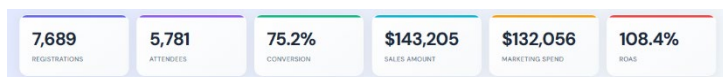
TOPIC	SALES AMOUNT	EXPENSES	ATTENDEES	CONVERSION %	QUALITY RATING
Creative Writing	21 944	13 573	562	13.88%	4.02
Ecology	21 417	24 204	740	9.32%	4.30
Graphic Design Principles	20 568	18 483	622	10.09%	4.22
Intro to Programming	20 439	18 973	745	7.62%	3.68
Business Management	17 544	22 467	982	7.29%	3.95
Digital Marketing	11 147	6 185	176	18.76%	4.05
Web Development	10 898	14 856	637	6.75%	4.00
Intro to Data Science	9 563	6 950	513	6.62%	3.88
Artificial Intelligence	7 175	5 931	531	4.34%	3.93
Introduction to Psychology	2 542	3 904	293	3.07%	4.80

The dashboard had 4 blocks but the table occupies 2 of them. I decided to abandon the trend chart since it doesn't show a clear trend. We could have changed the dashboard grid, but let's not get distracted by that for now.

## KPI Cards

Let's think about the final metrics. These are the first things the user sees; they should immediately reveal the topic. 3-4 numbers might not be enough, but overloading with up to 10 values isn't good either, as the meaning of a key indicator is lost. **Six** is the optimal number, but even then you need to be careful — what matters isn't the formal count of cards, but the number of meaningful content blocks.

In my version, Claude made 6 cards. But in terms of meaning, there are only two blocks here: one about the sales funnel (and this data is already in the chart), and the second – about the financial result and ROI.



To be clear, a card may repeat data from charts, but only partially, highlighting the most important elements. One card can bring together two or even three indicators.

- **Absolute (gross) value** – sales, headcount, units.
- **Indexed value** (with good/bad thresholds, as we did for the table). The most popular is sales plan fulfillment or expense limit adherence.

At this stage you can switch the AI from designer/layout mode to analyst/consultant mode. Let it suggest options following this logic, and you choose the ones suitable for your task. If you're not a manager yet and "don't know what's going on in the boss's head," then a dialogue with the neural network through the dashboard lens will help you formulate constructive proposals.

### Focus on money and profitability:

1. **Revenue** → ROI (%)
2. **Profit** → Margin (%)
3. **CAC** (customer acquisition cost) → LTV/CAC
4. **Average check** → Number of deals

### ✅ Option 3: Balanced (6 cards) ⭐ RECOMMENDED

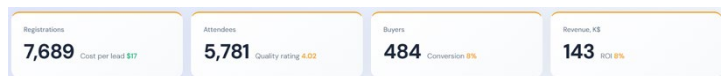
Full business picture:

1. **Registrations** → Cost per registration (\$)
2. **Attendees** → Quality rating (⭐)
3. **Buyers** → Conversion to purchase (%)
4. **Revenue, K\$** → ROI (%)
5. **CAC** (customer acquisition cost) → Average check (\$)
6. **Profit** → Margin (%)



In the end I got a set of 4 "absolute-indexed" indicator pairs:

- Number of registrations – Average lead cost
- Number of attendees – Average quality score
- Number of buyers – Conversion rate
- Revenue – ROI



To the indexed indicators I added conditional formatting, for example, considering a lead cost within \$20 as good. But this isn't necessary for every card; the additional

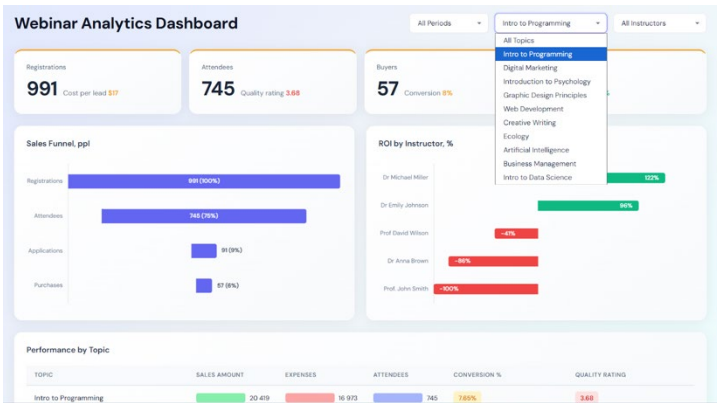
value can remain neutral, without "good/bad" thresholds.

## Filters and Interactivity

The final step – to bring it all together and add filters. After all, we need not just a slide, but the ability to view details by periods, departments, projects, branches. Or in other important dimensions you'll define in dialogue with the AI. Filters can be of several types, and the AI knows how to make them:

- Drop-down lists
- Slicers, similar to tiles with buttons, where you see selection options
- Sliders for number ranges or dates

You shouldn't overload the dashboard with too many filters; it's not a pivot table. In our case there are only three: by period, webinar topic, and lecturer. For this, you don't necessarily need a whole panel on the left; I placed them at the top, in the title area.



For now, this dashboard is only available inside the AI chat, but you probably didn't create it just for personal use. Let's figure out how to share it with colleagues. In Claude, you can create a link immediately, but ChatGPT needs instructions:

Generate an HTML page with an interactive dashboard for an Excel file. Use Plotly.js and SheetJS (XLSX). All data is processed on the client side, without a server. Add a form for file upload.

As a result you'll get a self-contained file that opens in any browser. You or another person you send the HTML template to will be able to upload fresh data and update the report.

Of course, such dashboards won't replace a full-fledged BI system. There are already separate AI-first services like Polymer, Bricks, and others. In them, you can implement the full data management cycle, access rights, simply by sending commands in a chat, without programming.

This topic would require a separate book, and most importantly – practice. If, when working with your own data, you can't get a beautiful dashboard in three prompts – don't despair. Join [the waitlist](#) to be the first to hear about our upcoming AI courses for business and career growth. You'll get hands-on practice with all document formats and explore many other topics designed to help you communicate effectively with neural networks.

# Business Infographics

Infographics are about conveying information through graphics. By combining text and illustrations, they create data visualizations, imagery, and visual metaphors that make information easier to perceive and absorb.

Numerous studies confirm that for visual communication – both with external and internal audiences – this works for comprehension and perception of material, and ultimately for everything that can be called selling an idea.

This genre became popular in the 2010s, along with the trend for data analytics and visualization. Gradually, any picture with graphics and icons started being called an infographic on social media. But we are talking about business applications, where you need to describe processes, connections, or create comprehensive content where charts, text, and diagrams are united by a visual image, a common idea.

This has always been a complex creative process, requiring professional design and understanding of the business context. Whole teams were involved – designers, analysts, editors. Only large companies or federal media could afford such departments. Now neural networks have advanced in creating visual

content. There are examples of truly high-quality designer infographics. But, as with other formats, the first attempts can be unsuccessful. In this part we'll break down what to expect from AI.

There are dedicated tools that were initially designed for image generation and later evolved to create infographics (or something resembling them). Here are the most popular ones:

- \* Midjourney
- \* Dall-e
- \* Ideogram
- \* Recraft

There are services that were popular for creating graphics and presentations 10 years ago, now enhanced with AI modules. In them you can now "dictate" an idea and develop it from a blank slate:

- \* Canva
- \* Infogram
- \* Visme
- \* Piktochart

If you ask your "everyday AI chat" for an infographic, it will certainly take on the job, but it's not a given it will succeed. In the role of a novice SMM manager, it might simply break the text into bullet points and add emojis, later calling it a technical specification. Claude in the role of an analyst might make you a technocratic diagram that's harder to read than the original text. Let's see what happens.

### **Case. Procurement Approval Process Diagram**

Descriptions of corporate regulations and business processes can span dozens of pages. On one hand, details are important, and nothing can be omitted; on the other – the larger the document, the fewer people will read it. In this example we are already using a simplified version of the process:

1. **Request Initiation** – The requesting department creates a procurement request (goods, equipment, services) with a description: specifications, quantity, deadline, justification of need, expense item code.
2. **Needs Approval** (by line managers) – The department head checks: is this procurement really needed, are there alternatives, isn't it part of an already approved plan. If rejected – request "rejected" or returned for revision.

3. **Budget Check / Financial Approval** – The finance department checks if funds are available in the budget line, correctness of the amount. If the amount exceeds limits – the request is returned or requires additional approval.

4. **Check of Procurement Norms / Internal Limits** – The procurement department (or process manager) checks: does the amount and conditions comply with internal limits for direct procurement, or is a tender, competition, request for quotations required.

5. **Launch of Vendor Selection Procedure** – The procurement method is determined: direct order, request for quotations, tender. Requests or invitations to vendors are prepared (three options if the method is "request for quotations"). Receipt of commercial proposals.

6. **Evaluation and Comparison of Proposals** – Collection and verification of proposals: price, terms, quality, reputation. Points or rankings are assigned. The best proposal is identified.

7. **Preparation of Final Proposal / Recommendation** – The procurement department or requester prepares a comparative report with a recommendation: winner and justification, potential risks.

8. **Approval of Final Proposal** – The department head and CFO approve the selected vendor, amount, conditions. If necessary – the

legal department reviews contract terms (if the amount is large or forms are non-standard).

9. **Request Approval and Order** – The manager (director, managing director) approves the request. The procurement department sends the order or contract to the vendor.

10. **Registration, Documentation, and Execution Control** – The final version of the request, contract, and protocols are stored in the document management system. Control of delivery deadlines, acceptance, payment.

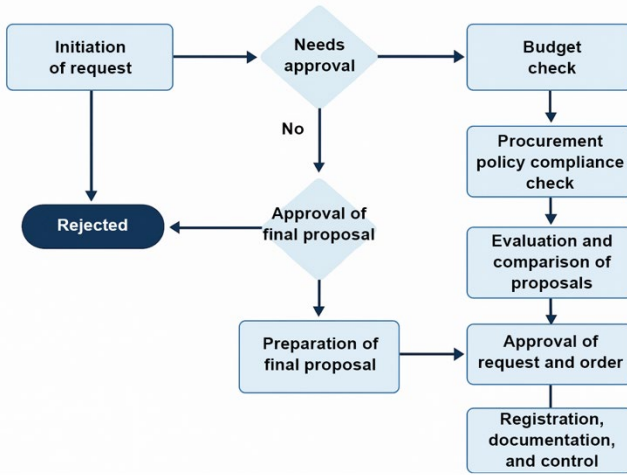
**Notes:** At steps 2, 3, 4 – requests can be returned for revision or rejected. During proposal evaluation, criteria may be revised or requests for revisions from vendors may be made. If there are legal objections – return to stage 8 or 7.

The prompt for now will be straightforward – "draw a business process diagram." Then we will refine it. And we'll use the "casting" method for neuro-infographers. Who handles the test task better, and who is worth spending more time communicating with. Let's start with Ideogram and ChatGPT.



Here's a version from the Ideogram service, which is tailored for creating pictures. It looks more like an illustration for a website, a cover for a procurement section. The metaphor is clear – office workers on a video wall are looking at a process diagram, or rather, at some image. And the language is strange. Now if only we could see the actual diagram, but alas – other generation attempts continue to yield only "imagery on the theme," not work with meaning.

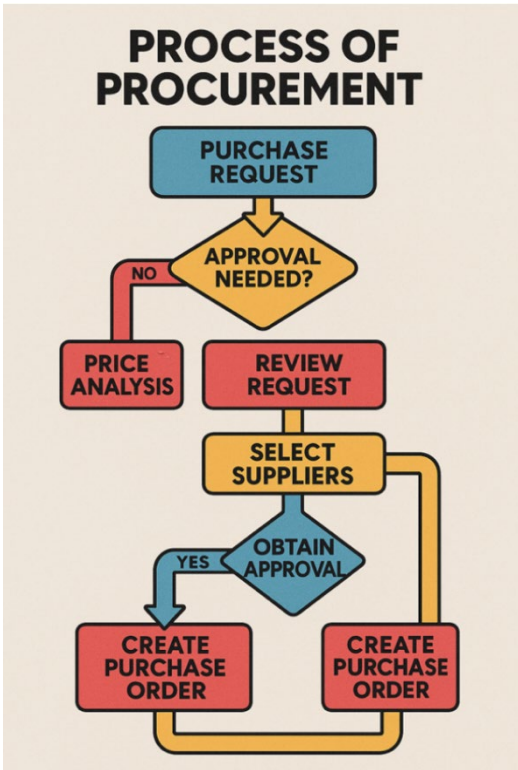
Conclusion – go to more experienced and business-oriented neural networks. Those trained not only on pictures but on documents. Let's start with the one with the largest training base – that's still ChatGPT.



Not bad for a first attempt: the diagram looks clear, concise, and finally, the full text in English. But if you look closely the logic is still broken: "Needs Approval – No – Preparation of Final Proposal". Let's say I won't pay attention to the logical inconsistencies for now and try to get some designer ideas from the AI:

I want to make a poster in a vertical A2 format to hang on the wall. And make it more vivid, infographic-like, add designer elements.

This variant still has a problem with logic: it no longer matches the original description.



I made another attempt to achieve a result:

Bring back all points from the original process that were in the very first version of the diagram. Besides the titles in the blocks, add a stage description, formulate them more concisely, as bullet points.

It got better, but still with logical inaccuracies. I asked the AI to compile a checklist for itself for verification. I got a two-page prompt, though when generating an image based on the new detailed requirements,



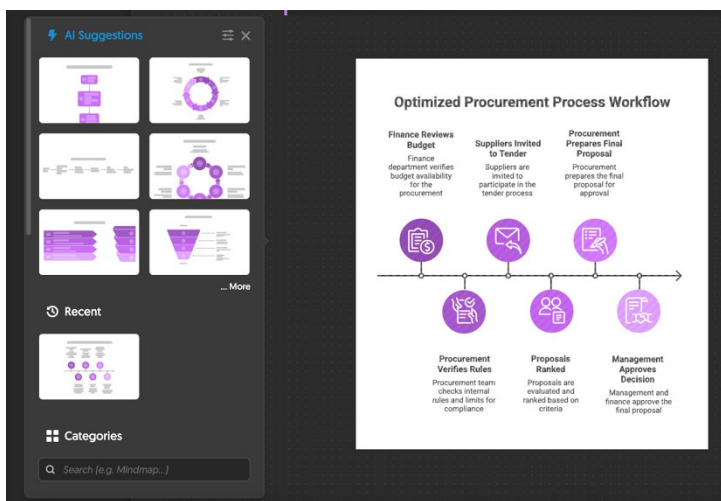
ChatGPT hung once and then gave an error. I started a dialogue in a new chat, and in the end, got a similar result in both.

If you had no luck with ChatGPT, I recommend trying Copilot from Microsoft. Working with office documents is in its DNA.

Overall, it turns out that in the role of an artist, the AI poorly monitors logic, simplifies, and confuses connections. Its "business drawings" can be taken as references and then given to a person for refinement. It's a pity it outputs images where you can't rewrite text or move icons. But specialized AI services can do that.

# Napkin.ai – Schematics You Can Edit

I like this service, its minimalist business style and ease of use. In principle, it resembles SmartArt in PowerPoint: there are schematic templates that adapt to your text. It looks incredibly simple: a blank canvas where you write text or paste ready-made text. You highlight lists, paragraphs, or all text, and then the magic happens:

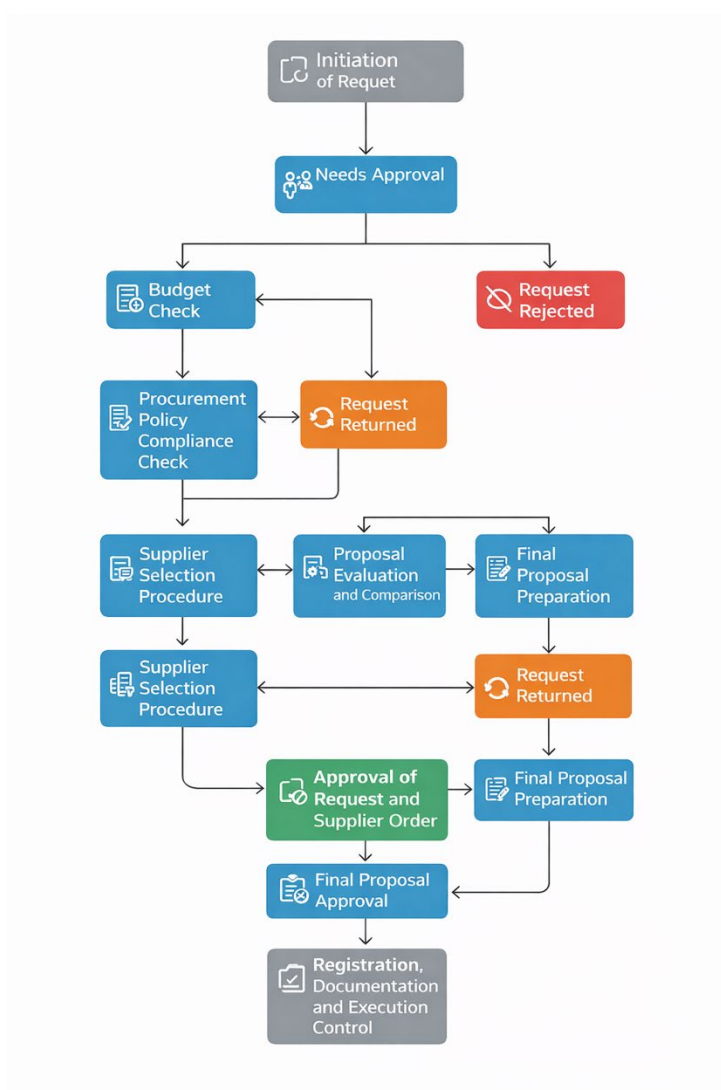


A panel appears on the left with different schematic options, and the number is unlimited! It will generate schemes until you find the right one (though at some point they'll start repeating). After deciding on the structure, you choose a style from pencil-like to bright contrasting. Then you get a vector diagram on the canvas and can edit it: change text, shapes, icons.

Remember, I wanted to make a poster in vertical A2 format to hang on the wall. Here's what suited me from the ready-made options.



The only thing I missed was the connections between stages, as there are process branches. For this I found a classic flowchart variant with decent icons, moved the blocks a bit, adjusted the connections, and here's what I got:



It reminded me of working with diagrams in MS Visio. Of course, Napkin won't replace professional tools and doesn't conform to BPMN notation. Its value is different – it combines balance, fairly diverse visual ideas, yet suitable in meaning. Honestly, I made the diagrams for this book in Napkin too. Not all of them turned out creative, but you can't make a carousel out of step-by-step lists.

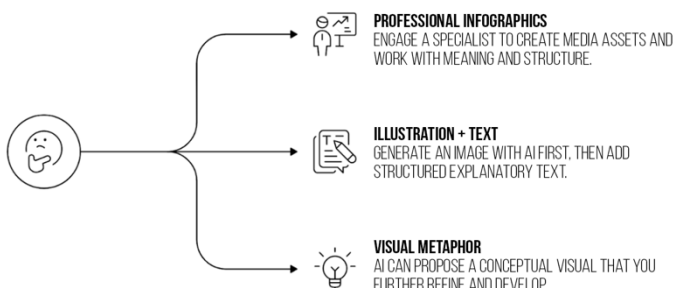
## Approach to Creating Infographics with AI

Step Zero – Make sure you really need an infographic, what task you want to solve with it:

- **Create a media project**, attract the attention of a wide audience with creative data visualization. In this case don't daydream and hire a person right away. Possibly, this specialist will use AI in the rendering process, but primarily, work with meaning is required: revealing statistics, showing interconnections, explaining complex data.
- **Create an illustration.** In broad circles, any combination of numbers, text, and images is considered an infographic. These are pictures

for social media, advertising banners, presentations. If you need a combination of a chart, an icon, and a text block, it's more reliable to make them separately. You've already seen – the neural network can't handle the roles of both artist and editor in one prompt. It's much easier to get creative illustrations and then add text in presentations or apps like Canva.

- **Come up with an image.** If you want to synthesize additional meaning from numbers and text, make an illustration that will facilitate the perception of a complex document or presentation, then you can try the neural network. There's a chance it might suggest a good idea you wouldn't have thought of yourself.

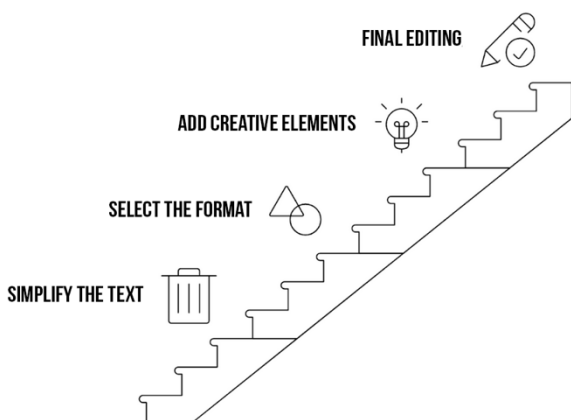


If you haven't changed your mind, the following steps await you:

1. **Dry out the text.** The AI in the process of generating a diagram, might suggest a structure with headings and short descriptions. But often these descriptions simply rephrase the heading without revealing the meaning. You need to carefully read each block and replace "fluff" with the essence – add keywords conveying meaning. Sometimes it's better to approve the text in a separate iteration.
2. **Choose the form.** This could be a cycle, a sequential process, or a map. Try different options and show them to a focus group to choose the best one.
3. **Add creativity and designer details.** Ask the neural network to make it brighter, more original, add elements that can breathe life into the diagram. Some ideas will be inappropriate, but eventually, a gem appears that you wouldn't have reached yourself.
4. **Final editing.** Look at everything as a whole: structure, texts, visual balance. Perhaps at this stage, you'll realize it's easier to assemble the final version in PowerPoint by hand than to

make the neural network correct mistakes over and over.

As you can see, I went through this process and at the editing stage I realized it's easier to split the diagram into two parts than to try to combine step zero, containing three options, with the subsequent four. Don't overcomplicate things, torture the neural network, or torture yourself checking its errors.



# Conclusion

Congratulations! You've reached the end of this extensive chapter, and we are getting closer to the final part of the book. I hope you have the opportunity to try applying the approaches I've told you about in practice. If not, it's all ahead of you!

Let's summarize: we considered visualization from the perspective of business tasks, based on data or texts. I am confident that in the future, humans will remain involved in this process. Sooner or later neural networks will learn to write without errors on diagrams and charts, but humans will be responsible for working with meaning and checking the result.

- **Charts.** To select by meaning and build a single chart based on data – AI already does this almost like a human.
- **Dashboards.** To assemble everything together, understand which metrics are important and which charts take up space without helping decision-making – that's a task for a human. But technically AI handles it well, indeed replacing a BI developer in simple projects.
- **Infographics.** In image generation mode AI handles this task poorly but AI-first services are actively developing; in them, you can get an idea and then edit the text, icons, and everything else.

## Chapter 7 Summary

Visualization for business tasks does not yet yield repeatable results. Careful human involvement is required at every stage.

### Dashboard Creation Steps

1. Define metrics and visualizations (AI will suggest)
2. Approve the dashboard structure
3. Configure charts, color coding
4. Configure KPI cards
5. Add filters
6. Create an HTML version

Best handled by: **Claude**

### Infographic Creation Steps

1. Dry out the text
2. Choose the form
3. Add creativity
4. Final editing

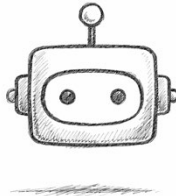
Best handled by: **Napkin**

**Development Trend:** Moving from image generation to creating editable objects, providing collaborative access.

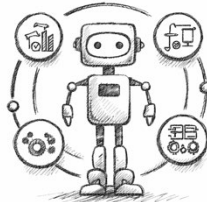
# Chapter 8. Business Process Automation



AI ASSISTANT



AI AGENT



RPA

In this chapter, we will try to bring together everything discussed in the previous chapters. Because above we broke down the elements of business processes; now we will talk about automating them as a whole. We'll look at what automation is in principle, what terms and tools are associated with it, and how to understand all this from a business perspective.

The goal of this chapter is to help you understand the basics. Every year thousands of new products and startups related to AI and automation emerge – dozens, hundreds of solutions for a wide variety of tasks.

Everyone promises to "do everything turnkey": automate processes, save time, reduce workload. Just choose our platform, and everything will work. But that's usually not the case. And often such promises are, in essence, a myth or soft marketing.

Therefore, my task is to give you fundamental concepts and systemic reference points. So you can answer the questions: What can truly be automated today? Which tools do what? Where can you get a ready-made solution, and where will you have to finish things yourself?

The chapter is structured as follows:

- In the first part, we will break down the terminology: what bots, agents, assistants are and how they differ. This will help you speak the same language as vendors and understand what is being discussed.
- In the second part, there will be a tool overview – examples of solutions from different classes and platforms. Nothing promotional, just specific examples so you see the whole picture.
- In the third part, we will analyse a practical example to understand how this works in reality.

Business process automation is the most extensive topic for applying artificial intelligence in business; not just a chapter, but a book would be too little. But I will still try to give you a general frame of reference because it is precisely in this topic where the largest field exists for inflated expectations and fascination with AI's capabilities.

Since effective process automation yields the greatest results, this is precisely where you can get the maximum synergistic effect from orchestrating people and machines. You train and adapt the machine, then it trains and controls employees, and the business only benefits from this.

# Terminology

## Robots and Bots

**Bots** are programs that operate according to pre-defined rules and interact with users through a messenger interface, web application, or corporate platform. They receive input events (messages, button clicks, commands) and return pre-determined responses or perform actions via API.

A classic bot does not make decisions on its own – it follows the logic embedded by the developer. For example, if a user clicks the "Order report" button, the bot calls a specific HTTP request and sends the result back. Therefore, bots are **deterministic**: with the same input, the result is always the same.

They are especially useful where a process can be expressed as a script with clear steps: request → check → action → response.

Typical examples:

- **Helpdesk systems.** The bot processes FAQ requests based on pre-written keywords or templates.

- **Forms and applications.** The user answers questions, and the bot forms structured data and transfers it to a CRM or database.
- **Notifications and reminders.** The bot reacts to events in a corporate system (a new lead, task, report) and notifies employees.
- **Status checks.** The user requests the status of an order, delivery, or booking, and the bot retrieves data from an external API and displays it in the chat.

## How a Bot Works

The bot's architecture is built around events and callbacks.

1. **Event Source.** The platform (e.g., Telegram, Slack, Teams) receives updates – this could be user input, new data arriving in the corporate system, etc. – and passes these updates to the bot.
2. **Dispatching.** An internal handler routes the event by type – text, button, command.
3. **Logic.** The corresponding function is called (e.g., `handle_order()` or `get_status()`), where the behavior is described.

4. **Response to User.** The platform sends a message, image, file, or link via an API method.

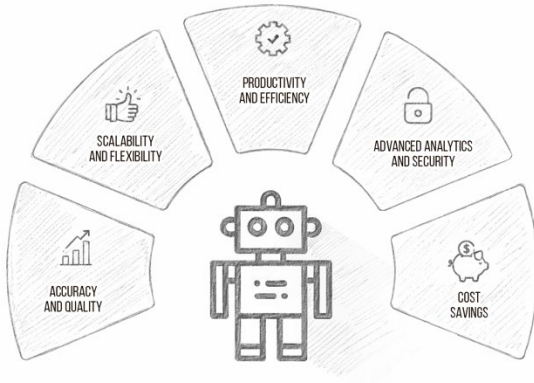
**Robots** in the context of automation are software performers that mimic human actions in application interfaces. They do not communicate with the user directly, like a chatbot, but perform operations within an information system, copying human behavior: clicks, data entry, reading tables, switching between windows.

This approach is called **Robotic Process Automation (RPA)**.

## Robotic Process Automation (RPA)

The main idea of RPA is automating routine tasks at the user interface level without changing the source code. Its tasks are to perform monotonous work that takes a lot of human time but does not require a creative approach, for example, data entry, document processing, report generation, etc.

If a company has dozens of legacy systems (CRM, ERP, accounting, Excel files) that are difficult or expensive to integrate, RPA allows "stitching" them together through the visual interface without touching the internal APIs.



## How RPA Works

The robot "sees" the interface and performs a sequence of actions as a human would:

1. **Perception.** The robot captures the screen image, form fields, menu elements, and recognizes them (via OCR, templates, DOM structures). For example, it can recognize a table in Excel, a form in SAP, or an input field in a browser.
2. **Processing.** The action logic is described as a script or flowchart: "open application → enter login and password → find document → copy data → paste into CRM form." The algorithm is executed step by step, often without human involvement.

3. **Action**. The robot clicks, types, opens, copies – that is, reproduces physical gestures but in a digital environment. It doesn't just call an API; it mimics interface interaction, which makes it universal but less flexible compared to code-based integration.

## Types of RPA Robots

1. **Attended** (Human-controlled). Launched manually. They assist an employee by performing part of the routine steps – for example, processing a form or filling out a template.
2. **Unattended** (Fully Autonomous). Operate on a schedule or based on triggers. Example: a robot checks for new orders every hour, downloads files from email, populates Excel, and enters data into ERP.
3. **Hybrid**. Combine both modes: some scripts are launched by the user, others automatically. Such solutions are more common in corporate RPA platforms where robots are integrated into a unified orchestration system.

# RPA Architecture

RPA platforms (UiPath, Blue Prism, Automation Anywhere, Power Automate Desktop) include several levels:

1. **Designer** (Script Designer). The interface where a business analyst visually describes steps: actions with windows, clicks, data entry, loops, conditions.
2. **Robot** (Performer). The agent program installed on a computer or server that executes the script.
3. **Orchestrator** (Coordinator). The central control panel for robots: who does what, when, with what success, under what load. Schedules, access rights, monitoring, and license management are set here.
4. **Recorder and Computer Vision**. Modules that record user actions and turn them into scripts. Using CV and OCR, robots can recognize visual elements even if their on-screen coordinates have changed.

## Examples of Practical Use

- **Bank Back Office:** A robot checks loan applications, opens emails, downloads attachments, enters data into CRM, and initiates scoring.
- **Accounting:** A robot takes bank statements, reconciles amounts with ERP, and generates a report.
- **HR:** A robot collects CVs from portals, enters them into the system, and sends a notification to the recruiter.
- **Legal Department:** A robot monitors the government procurement or arbitration cases website and updates tables with results.

## Limitations and Vulnerabilities of RPA

1. **Dependence on the Interface.** Any change in the design or page structure can "break" the script – the robot won't recognize a button or field.
2. **Lack of Context.** The robot does not "understand" the meaning of data – it only follows instructions. Any unexpected error (a new window, a different message) requires manual script updates.
3. **Scaling Complexity.** With a large number of scripts, it is difficult to manage versions, updates, and security.
4. **High Implementation Cost in Large Systems.** Despite visual tools, corporate RPA requires licenses, infrastructure, and support.

## Agents and Assistants

Now – about roboticization. Most often, this word refers to corporate scenarios: when a "virtual employee" appears in a company that works between applications. It opens interfaces, clicks buttons, downloads documents, transfers data. Essentially, it repeats the actions of a living person.

**Important:** These are not industrial robots or "iron arms." We are talking about virtual roboticization, i.e., simulating user actions in digital systems. Not even exactly an employee, but precisely an imitation of their behavior in interfaces.

In the book I compare different options: I analyze bot-robots, agents, and assistants. Based on clear criteria – so that it's interesting to read myself and easy to grasp the difference. I've heard about roboticization for a long time and from everywhere, but only while working on this material did I finally understand its real essence. And now I want to share this understanding with you.

I define the terms and immediately offer the criteria by which they can be distinguished from each other. What exactly makes a robot a robot? At what point does an assistant stop being just a helper and start behaving like

an agent? And where is the line between these approaches?

The same goes for roboticization. It has obvious pluses, but also minuses that are usually mentioned in passing. For example, cost: in reality, such automation often turns out to be noticeably more expensive than expected at the start. Especially in corporate implementations where you need to maintain scripts, interfaces, and process stability.

So who are AI agents and assistants? They are comprehensive tools for working with AI, and their main difference from each other lies in their type of activity. Agents are **reactive**, meaning they do not make decisions themselves; they need a command from a human. Assistants are **proactive**, meaning they possess a certain degree of autonomy and can make their own decisions. Let's examine them in detail.

## Assistants



As I said, they are **reactive**, meaning they respond to human actions who initiate the interaction, for example, by asking a question to which they will give an answer. They are useful for accelerating a human's work in tasks they are already performing.

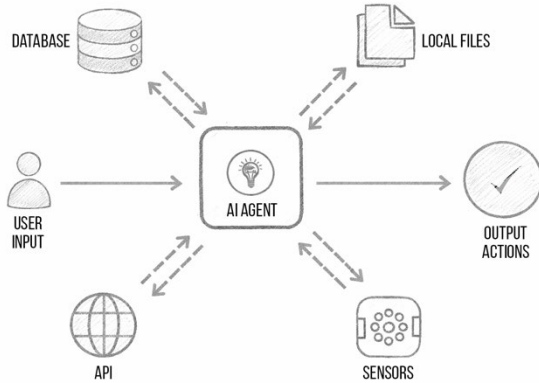
Despite their high utility, the productivity gain they offer is more limited compared to a fully autonomous agent. AI assistants can be useful in a wide variety of tasks, for example:

- **Composing emails.** Tools like ChatGPT help prepare a correct response to a client,

considering previous correspondence and the desired tone, without the employee needing to manually re-read the entire message history.

- **Conducting research.** An assistant can quickly search for information online and select key facts to answer a posed question.
- **Working with extensive documentation.** Thanks to semantic search, they find relevant information by meaning, not just by keyword matching.
- **Writing code.** Specialized coding assistants significantly speed up the development process and increase programmer productivity.

# Agents



An **AI agent** is an application based on artificial intelligence that is capable of acting **proactively**, i.e., autonomously, making decisions, and performing tasks with minimal human intervention. Thanks to the use of AI methods, such an agent can analyze data, learn from experience, and adapt to changes in the external environment, increasing the efficiency and accuracy of completing assigned tasks. Agents can perform actions without direct instructions. Because of this, AI agents are useful for multi-stage workflows, especially those that interact with other systems.

As of 2025 both assistants and agents are used for process automation and decision support – examples

can be seen in solutions like IBM Watson or self-driving cars. The difference is that an assistant helps a specialist by carrying out assignments, whereas an agent can independently take on some of the responsibilities and act more autonomously.

Agents have other important characteristics. They can be launched from other systems using an event-driven model. They can also interact with other data sources (in some cases experimenting with creating and testing their own APIs) to perform actions or obtain additional information for processing, for example, as part of a RAG pipeline (Retrieval-Augmented Generation). AI first finds the necessary information and then answers based on it. Comparing agents with robots makes it easier to explain their working principle.

Their activity can be divided into three stages:

1. **Perception.** The agent receives information about the environment through sensors. These sensors can be physical (like in self-driving cars) or virtual – for example, when receiving data from an API.
2. **Processing.** Using incoming real-time data, the agent analyzes the situation and chooses the most appropriate solution. Unlike traditional systems that have no connection to external data

sources, this approach allows for considering context.

3. **Action.** After making a decision the agent affects the surrounding environment using "actuators." This can be a physical action (turning a car's steering wheel) or a virtual one – for example, changing a record in a database via API. When the agent changes the environment, it again uses sensors to record the result and consider it when choosing the next step.

Although the agents themselves (and the neural networks they run on) do not learn during task execution, their work allows for data collection for subsequent model training and improvement. This process over time makes agents more effective and enables their use in more diverse conditions with fewer limitations. Today, in most cases, agents interact with the world using virtual sensors and actuators. However as their reliability grows we can expect them to start managing an increasing number of physical systems.

# External Agent Functions

External agent functions are tools and resources that expand the agent's capabilities. The instructions must describe which tools the agent can operate with – for example, vector databases for semantic search or APIs for obtaining current data. The agent independently chooses which tool to use at each stage, based on the nature of the task, while operating within predefined guardrails. **AI Guardrails** are restrictions and security policies that ensure the agent operates within ethical and functional boundaries.



For example, you can set an instruction to maintain a unified tone of voice (ToV) for the brand so that the agent does not slip into an inappropriate communication style (e.g., making unfounded promises or using slang inappropriate for the company's context). One guardrail is clearly not enough. View them as a multi-layered

protection system, where each guardrail is responsible for a certain class of potential risks.

It is also important to include error handling logic in the prompt: how to detect failures, when and how to retry operations, and in which cases it is necessary to escalate the problem. For example, the instruction might be: *"If an external API causes a timeout, record the error, wait 5 seconds, and retry up to three times. If unsuccessful, notify the user."*

**Memory** is a storage of previous interactions, facts, and results of tool usage that the agent can access during its work. You can set precise instructions on what information to save and at what moments. This could be user preferences, tool operation results, extracted entities (e.g., dates or locations), important milestones, etc. You can also manage how the agent compresses and cleans memory to avoid overload: for example, by summarizing past interactions, deleting outdated data, or prioritizing critical information.

# Tools

Here I will talk about the specific tools we can use to create AI assistants of different levels. Let's consider three levels: bot builders, visual integration platforms, and RPA platforms.

## Simple Bot Builders

We're talking about the most basic level: bot builders that react to incoming events. These are the "simple" or, to be honest, rather "dumb" bots. They don't reason, plan, or make complex decisions – they simply do what they're told when a specific event occurs.

For example: an email arrived → launch a script. A new row appeared in a table → pass the data further. A button was clicked → execute a chain of actions. All logic here is built on the "if → then" principle.

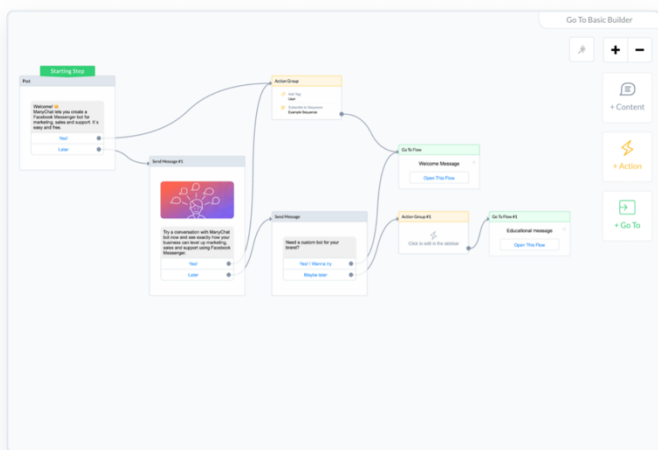
Usually, such tools look like visual diagrams. Blocks, arrows, triggers, actions – you literally assemble the process with a mouse. That's why there are pictures next to this section in the book: to immediately make it clear what this even looks like and why it's not about code in the classical sense.

Below are several popular tools from this category. This is not a rating or a crash test, but a short navigation guide: so you can orient yourself on what is suitable for what and whether it makes sense to pay for it.

The most accessible way: choose a template, write an instruction – and the bot in the messenger is ready.

## **ManyChat**

ManyChat is a US-based platform for building chatbots and automated funnels. The product appeared in the mid-2010s and was originally built as a tool for marketing, sales, and audience engagement through messengers and social platforms. Its main focus is Instagram, Facebook Messenger, and WhatsApp, with integrations into CRM systems, email tools, and other marketing services. ManyChat was one of the first no-code bot builders that became widely adopted by marketers, creators, and online businesses in the US. It is strongly oriented not toward developers, but toward marketers and growth teams, and is often used as a “control center” for funnels: bot flows + broadcasts + lead capture + basic CRM logic.



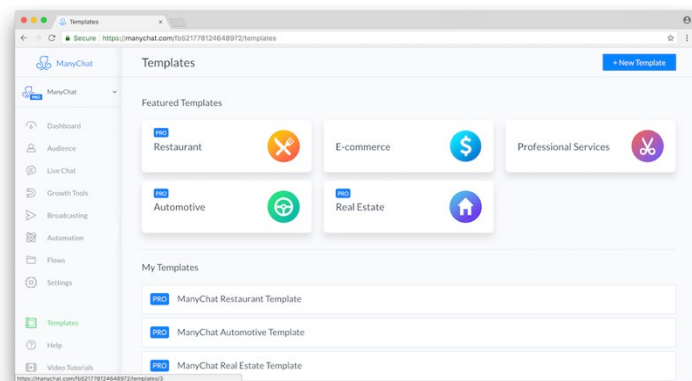
Source: <https://manychat.com/>

On the screenshot above, you can see ManyChat’s visual flow builder. This is a typical event-driven flow with branching logic: a user enters the flow → sees a welcome message → clicks a button → triggers an action group (tags, subscriptions) → and is redirected to another flow.

The logic here is not linear, but conditional: user choices (“Yes”, “Later”) define what happens next. Messages, actions, and transitions between flows are visually connected, making it clear how user input controls the path through the automation. ManyChat covers tasks related to event-based automation and user communication:

- sales and warm-up funnels;
- lead capture and qualification;
- broadcasts and timed message sequences;
- collecting user data and passing it to external systems;
- basic CRM-style logic inside flows.

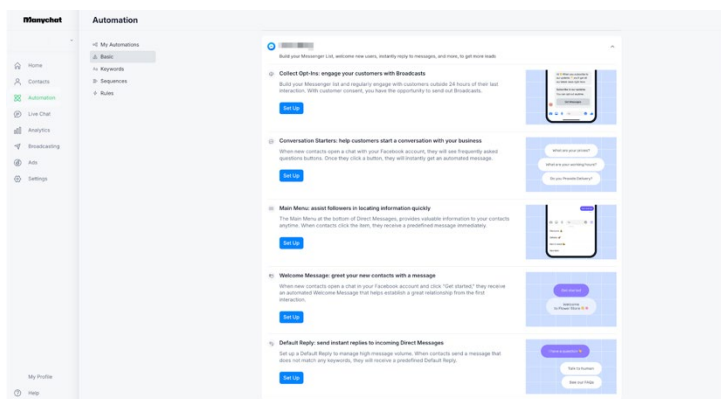
By nature, this is a classic event-driven bot: the user does something → the bot responds according to a predefined flow.



Source: <https://manychat.com/>

On this screenshot, you see ManyChat's template library — a categorized gallery of ready-made flows grouped by business type, such as e-commerce, restaurants, real estate, and professional services.

Here, ManyChat presents itself not as an empty builder, but as a system with pre-packaged business scenarios. You don't have to start from scratch: you can choose a template that matches your use case and adapt the flow to your task.



This screen shows ManyChat's core automation building blocks — the basic ways the platform reacts to user behavior. These are predefined entry points where communication starts.

At this level, automation is about structure. You decide when the bot should speak, where the conversation should begin, and what kind of interaction it should be. E.g. a user enters — they get a welcome message, or someone clicks a button — a conversation starts. If a

message doesn't match any rule — a default response is sent.

These automations define the boundaries of the system: how users enter the flow, how they are guided, and how predictable the interaction remains. Everything here is intentional and controlled. This is the foundation on which more complex logic and AI can be safely layered.

What ManyChat adds on top of AI:

- control over tone and role;
- connecting AI responses to flow logic;
- use inside messengers and social platforms;
- moderation, limits, and safety controls.

Even with AI enabled, ManyChat does not turn into an agent system, which we'll talk about later.

ManyChat does *not* have:

- independent planning,
- goal setting,
- long-term autonomous memory,
- actions outside predefined flows.

This is an AI-assisted layer inside an event-based bot — a practical and safe format for business.

## Tidio

It is a no-code bot builder designed primarily for live chat on websites and simple communication scenarios. Tidio is usually used as an on-site chat widget that talks to visitors while they browse a page, look at a product, or decide whether to leave their contact details. Functionally, this is very close to a messenger — the difference is that the conversation lives on the website, not in Telegram or WhatsApp.

The logic is still event-based. A visitor opens the chat, clicks a button, types a message, or triggers a rule by their behavior on the site — and the predefined scenario starts.

Inside Tidio conversations are built as visual flows. Messages are connected with arrows, and certain blocks change the direction of the dialogue depending on what the user does. For example, the chat can ask a question with two possible answers, and each choice leads the visitor down a different path. The structure stays simple and readable: you always see where the conversation starts, where it branches, and where it ends.



- This is not an agent or an assistant, but a rule-based, event-driven chat;
- Little flexibility for highly custom processes or advanced analytics.

Tidio works best when the task is straightforward. When you need to quickly set up a website chat that reacts to visitor actions and guides them through a predefined conversation. That's exactly what it's meant for — and in that context, paying for it makes sense.

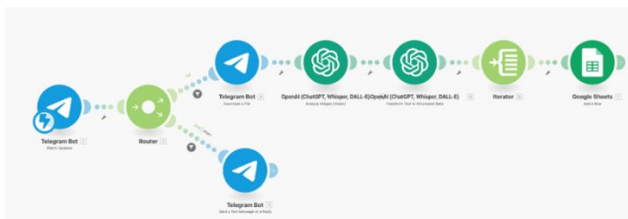
## Visual Integration Platforms

These are platforms where you can independently connect different services (CRM, databases, messengers) into a single automated process according to your scenario. Importantly, no programming is needed to create the action chain; ready-made nodes (blocks with predefined functions) are used.

### **Make**

This is one of the most well-known visual automation builders. It started as **Integromat** and was later renamed **Make**.

Here you can build entire workflows: a website lead goes into CRM → the client receives an email → data is sent to Google Docs.



**Example:** A request for a commercial proposal arrives by email. The system acts on its own according to a set scenario:

1. Reads the email.
2. Extracts the client's name and address.
3. Records the data in CRM.
4. Finds the CP (commercial proposal) template.
5. Generates the text.
6. Sends it to the manager for review.
7. Sends the client a beautiful email with the ready CP.

The entire chain – without human involvement.

**Pros:**

- Flexibility; you can build almost any chain without code;
- Rich capabilities for working with standard and even non-standard tasks;
- There are ready-made scenario templates.

**Cons:**

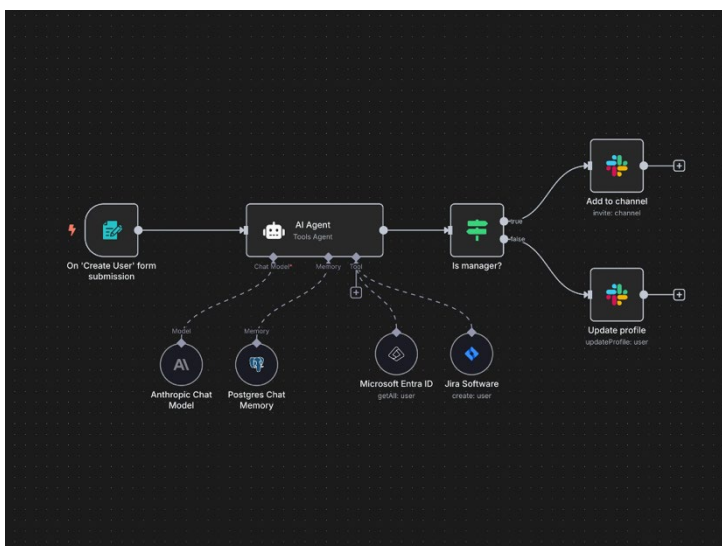
- Data passes through the internet – there's always a security question;
- Even with a template, you'll have to figure out the settings of each module yourself;
- If something breaks, you don't see the code; you need to dig into the interface and find where the "wire broke";
- For some business tasks, the functionality still won't be enough – there are nuances that such services don't cover.

## n8n

An open-source alternative to commercial builders. Often heard of as "Make, but for those who want more control."

You can deploy it yourself, or use the cloud version.

Suitable for those who want flexibility and aren't afraid of a slightly more technical interface. Often chosen by teams and developers for whom not being completely dependent on a SaaS platform is important.



(Source: [n8n.io](https://n8n.io))

How processes are structured schematically in n8n:

- Workflow — the main object. It's a visual diagram of nodes connected to each other.
- Trigger — the starting point of the process. Activated by an event: webhook, schedule, data change, incoming request.
- Nodes — individual steps (nodes) of the process. Each node performs one action: get data, transform, send further.
- Data flow — data is transferred between nodes in JSON format; each next step works with the result of the previous one.
- Conditions — logic branching. You can check values and direct the data flow along different paths.
- Loops — processing data arrays: n8n can iterate through lists and process elements sequentially.
- Custom logic — if necessary, you can insert code (JavaScript) directly into the process.
- Execution history — run history, where you can see what worked and where the process broke.

In general, the logic is: event → sequential chain of steps → data processing → result.

**Pros:**

- Very flexible logic compared to classic builders;
- Support for complex scenarios, branching, and loops;
- Open-source: can be self-hosted, independent of SaaS;
- Well-suited for integrations between different systems;
- Ability to write custom code directly in scenarios;
- Transparent work with data; you can see what is being passed where.

**Cons:**

- The entry threshold is higher than for Make or Zapier;
- The interface is less user-friendly for business users;
- Without a basic understanding of JSON and process logic, it quickly becomes complicated;
- Self-hosting requires support and administration;
- Debugging complex scenarios can take a lot of time.

Make and n8n became popular due to the balance between clarity, flexibility, and the number of integrations, and because they cover most basic automation scenarios without code.

Below I show a few more tools, briefly and without fluff, so you can compare approaches and understand which option better suits your task.

## **Zapier**

One of the oldest and most mass-market automation services. Made the very idea of "connecting everything to everything" popular. Very simple to start: choose a trigger, choose an action – it works. Good for simple scenarios and business users.

Limitation – the logic is quite linear, and complex chains quickly run into pricing limits.

Typical tasks that can be solved with Zapier:

- Transferring leads from website forms or Google Forms to CRM and a manager's email;
- Automatically creating tasks in Trello / Asana for new events (deal, email, application);

- Synchronizing data between Google Sheets and SaaS services (contacts, statuses, tags).

**Pros:**

- Very simple start, almost no training required;
- Huge number of ready-made integrations;
- Stable operation for simple scenarios.

**Cons:**

- Quickly becomes expensive with volume growth;
- Limited logic and weak branching;
- Difficult to scale non-standard processes.

### **Pabbly Connect**

Less well-known but often mentioned as a more affordable alternative to Zapier. Functionally similar: events, actions, connections between services. Suitable for small business processes when you want to automate routine but not overpay for the brand.

Typical tasks:

- Transferring data between forms, spreadsheets, and email services;
- Automatically adding users to mailing lists and lists;
- Simple business processes: applications → spreadsheet → notification.

**Pros:**

- More affordable cost compared to Zapier;
- Support for basic automations without code;
- Suitable for small teams and startups.

**Cons:**

- The interface is less convenient and clear;
- Fewer ready-made templates;
- Harder to debug chains when errors occur.

**Albato**

Works well with local services and popular CRMs. Often chosen for its clear interface and focus on business tasks, not technical details.

### Typical tasks:

- Transferring applications from forms and landing pages to CRM;
- Synchronizing data between CRM, spreadsheets, and email;
- Automating notifications to managers and clients.

### Pros:

- Clear interface without overload;
- Often more cost-effective than Zapier.
- Deep integration with local services that are critical for day-to-day operations

### Cons:

- Albedo was developed in Europe and its integration catalogue is less comprehensive for US-centric, or niche SaaS tools
- Limited logic for complex scenarios;
- Less flexibility compared to n8n.

# RPA Platforms

We've already talked about what RPA is (see the section *Robots and Bots*), so here's a brief reminder of the essence. RPA platforms (Robotic Process Automation) are needed to automate routine operations that a person performs manually in program interfaces: opens applications, switches between windows, copies data, clicks buttons, and fills out forms. Essentially, the robot repeats user actions step by step, strictly according to a given script.

Examples of practical use:

- **Bank back office:** A robot checks loan applications, opens email, downloads attachments, enters data into CRM, and initiates scoring.
- **Accounting:** A robot retrieves bank statements, reconciles amounts with ERP, and generates a report.
- **HR:** A robot collects resumes from portals, enters data into the system, and sends a notification to the recruiter.
- **Legal department:** A robot monitors government procurement or arbitration case websites and updates tables with results.

RPA platforms work at the interface level. The robot is shown where to click, which fields to fill, and in what order to perform actions. It sees the screen, recognizes interface elements, and repeats the script as if a human were sitting at the computer. This is the main difference from bots and automation builders we discussed earlier. Bots work through events and integrations – they don't need a screen; they get data directly from systems. RPA is used where integrations don't exist, are unstable, or are too expensive to develop.

RPA platforms are especially effective in large companies with legacy systems, complex regulations, and a high volume of manual work. Where processes already exist but are performed by people for hours, a robot can provide a quick effect without reworking the IT landscape. However, RPA also has pitfalls. Such robots are sensitive to interface changes: the screen updates – the script may break. Support and maintenance cost money, and complex processes turn into fragile chains.

Therefore, it makes sense to use RPA where you need to automate precisely "human" actions in systems. If a process can be handled through events and integrations, regular bots and automation builders are often simpler, cheaper, and more reliable.

## Top 3 Global RPA Platforms

### UiPath

The most massive RPA vendor in the world, went public in 2021. The platform has become the de facto standard for enterprise RPA. Founded in Romania in 2005, later headquarters in the USA.

Suitable for tasks:

- Mass automation of back-office processes;
- Finance, accounting, banking operations;
- Automation of legacy systems without APIs.

**Process:** Robots are assembled in a visual designer from steps and conditions. Supports work with interfaces, files, browsers, OCR. Robot management, schedules, control, and logs are handled through a centralized Orchestrator.

#### **Pros:**

- Very mature ecosystem and stability;
- Support for complex scenarios and scaling;
- Large community and documentation.

**Cons:**

- High license costs;
- Complexity of implementation without an experienced team;
- Overkill for simple tasks.

**Automation Anywhere**

One of the oldest and most widespread RPA platforms on the global market. Initially created for large companies and large-scale roboticization programs. They emphasize centralized robot management, security, and operation in a cloud architecture. Often used in an enterprise environment where control, access roles, and process manageability are important.

Suitable for tasks:

- Corporate regulatory processes;
- Automation of operations in finance and HR;
- Large-scale RPA programs in enterprise.

**Process:** Scenarios are assembled in a visual editor, using commands, conditions, and triggers. There is centralized management of robots, roles, and security.

**Pros:**

- Well-suited for large organizations;
- Cloud architecture;
- Developed access management tools.

**Cons:**

- Interface is less intuitive than UiPath;
- Requires training and methodology;
- Total cost of ownership.

**Blue Prism**

A system from the UK, one of the most "strict" RPA vendors, initially created for banks and financial organizations. Oriented towards strictly regulated processes. Therefore, it emphasizes reliability, reproducibility, and audit of robot actions. Suitable for scenarios where predictability and rule compliance are important.

**Processes** are modeled as business logic with clear steps and rules. Minimum "freedom," maximum control, audit, and reproducibility.

Suitable for tasks:

- Processes with strict regulations;
- Financial and compliance processes;

- Critical operations with high control.

**Pros:**

- Very high reliability;
- Good for regulated environments;
- Strict process architecture.

**Cons:**

- High entry barrier;
- Little flexibility;
- Slow to implement changes.

Now we have a general picture. We've broken down different automation entities and seen where each is appropriate. AI assistants, AI agents, and RPA solve different tasks and confusion between them usually arises when a tool is expected to do something it's not designed for.

To conclude the theoretical part, let's recall that:

- An Assistant is primarily a helper for a person: it suggests, analyzes, formulates, accelerates decision-making, but doesn't have its own goal.
- An Agent is a step further: it has a task, context, and the ability to act autonomously, choosing steps and adapting along the way.

- RPA is a completely different plane: it's the precise and reproducible execution of human actions in interfaces.

It's also important that these approaches do not necessarily compete with each other. They can be combined, but they shouldn't be mixed up in one's mind. RPA is effective where manual labor in systems without APIs or stable integrations needs to be replaced. AI assistants are good for supporting people – analysts, managers, lawyers, HR – for example, when meaning, text, interpretation, and recommendations are important. AI agents are appropriate in more complex scenarios requiring autonomy and decision chains. Sometimes assistants and agents can manage RPA robots, and sometimes they can do without them. The main thing is to understand what exactly you are automating: actions, thinking, or decision-making, and then choose tools accordingly, not the other way around.

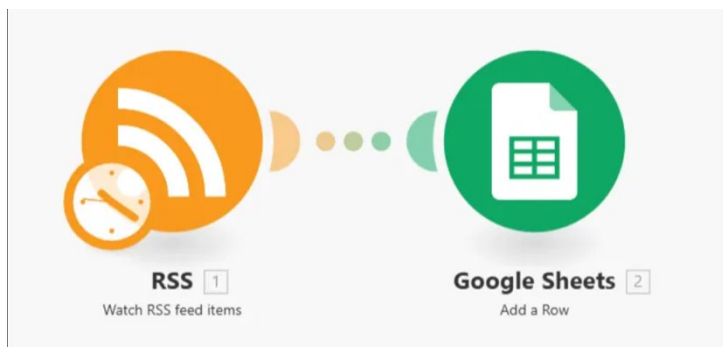
## Practice. Automation Template for Competitor Monitoring

Let's consider an example of how to automate competitor activity monitoring. Such a system is needed for decision-makers who don't want to rely on gut feelings. It's very difficult to manually track what competitors are doing, when and who launched a new product, who is strengthening their team, and similar things. As a result, decisions are made late or blindly.

Automated monitoring solves this problem. In short, it's a tool for businesses where reaction speed and awareness are important. This automation template does not replace analytics and strategy and does not provide ready-made solutions, but creates a foundation for them, optimizing time that could have been spent on manual labor.

The process describes a system for automatically monitoring competitor activity: publications, news, job postings, and changes in public sources. It includes data collection via RSS, Telegram channels, and website pages, automatic aggregation into Google Sheets, and formation of a unified event log. It also shows how to move from collected data to analysis and conclusions

using a neural network – without manual monitoring and constant source checking.



No exotic tech stack or rare tools. We're talking about services most have already encountered: an account on Make.com, a Google Sheet, and data sources like RSS feeds (Really Simple Syndication), Telegram channels, or competitor websites. This minimal set clearly illustrates an important point: the value lies in how they are interconnected and what problem they solve.

Create a Google Sheet with the following structure:

	A	B	C	D	E	F
1	Date	Source	Type	Title	Link	Description
2						
3						

What will we monitor?

We take the simplest and most obvious sources – the ones usually monitored manually.

- First, **RSS feeds**: Competitor news and blogs. Everything new is released and immediately comes to us without needing to visit the website each time.
- Competitor **Telegram channels**. A quick way to understand what they are talking about, launching, and emphasizing.
- Competitor **job posting sites**, both aggregators (like LinkedIn Jobs, Indeed or ZipRecruiter) and their own career pages. These show well who and for what tasks they are currently hiring, and thus, where the company is heading.

## Step 1. Setting up the Make scenario

Create a new scenario in Make.com.

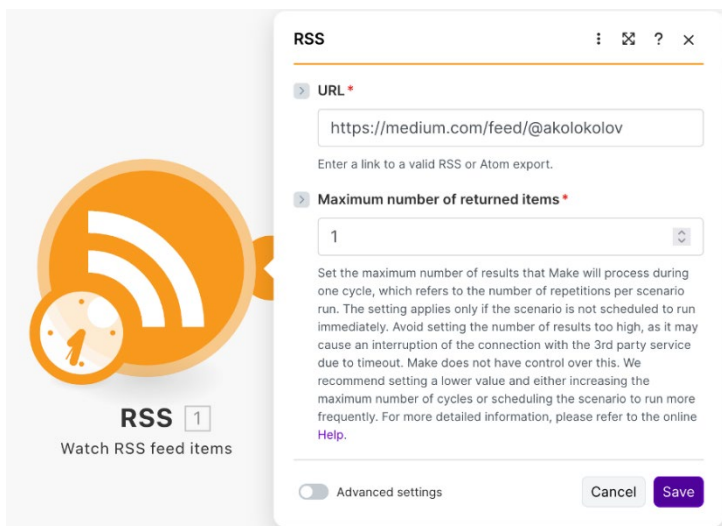
### **Module 1: RSS (for news or blog sites)**

The task here is simple: teach the system to notice changes for us.

1. Add a module: RSS → Watch RSS feed items;

2. Paste any RSS link, e.g.:  
`https://www.wired.com/rss` or  
`https://example.com/blog/rss`;

3. Set the update interval: every 15 minutes / once a day.



## Module 2: Google Sheets → Add Row

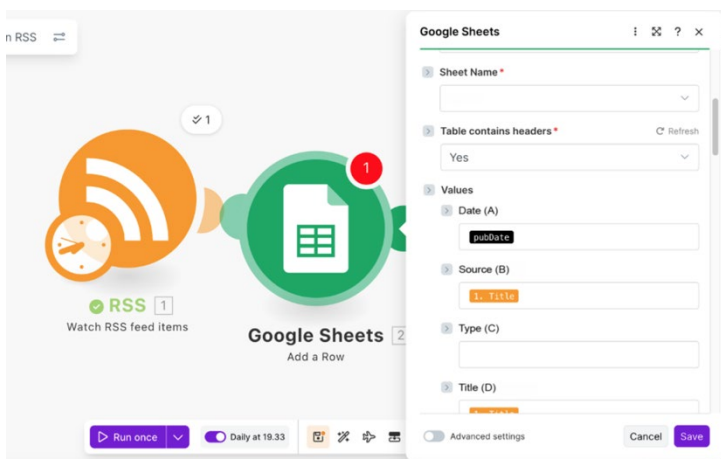
At this step, we get a collection point. Everything that was previously scattered across sites, channels, and pages starts neatly piling into one spreadsheet.

1. Add a module: Google Sheets → Add Row and select the spreadsheet and desired sheet.

2. Fill in:

- \* **Date:** `{{pubDate}}`
- \* **Source:** Channel name
- \* **Type:** News (or Job Posting)

- \* **Title:** `{{title}}`
- \* **Link:** `{{link}}`
- \* **Description:** `{{description}}` (if available)



That's it! Every new article/job posting will go into the spreadsheet. You can then do anything with this data: filter, sort, analyze. But the main thing is that information stops being chaotic and becomes manageable.

## Step 2. Adding Telegram Channels

Telegram itself doesn't support RSS, but there are free converters. Essentially, nothing new happens. We simply disguise Telegram as a regular news source.

After that, it starts behaving like a blog or website: new posts are published, the scenario notices them, and then everything works according to the familiar logic.

Here's a way to convert a channel into a system-understandable format:

### Via an RSS converter

1. Take a Telegram channel link, e.g.: `https://t.me/example_channel` and to form the link, use a service like `https://tg.i-c-a.su` or `https://rsshub.app/telegram/channel/example_channel`
2. Insert the resulting RSS link into Make just like a regular RSS (see step above).

#### RSS

```
- URL: https://tg.i-c-a.su/rss/breakingmash  
- Custom limit: https://tg.i-c-a.su/rss/breakingmash?limit=50  
Maximum posts - 100 posts  
- Pagination: https://tg.i-c-a.su/rss/breakingmash/2  
- One post: https://tg.i-c-a.su/rss/breakingmash/?id=10738&limit=1
```

## Step 3. Monitoring Job Postings

You can:

- Find RSS (if available), or
- Use Visualping.io or Distill.io to track changes on pages (if you don't want to code)

But for simplicity, you can:

1. Find a competitor's job posting page;
2. Connect Visualping – it will send an email upon change;
3. Email → Make → Spreadsheet (similar to the previous scenario).

Example sources:

Company Name	Monitoring	Tool
Competitor A	Newsfeed	RSS
Competitor B	Telegram Channel	RSS via RSSHub
Competitor C	Job Boards	Visualping + Email

## Step 4. Testing

The final step in our chain and the moment of truth.

- Wait for a publication in RSS / Telegram;
- Check if a record appears in the spreadsheet.

When you see the first record appearing by itself, it becomes clear that the chain is closed. The source published news, the scenario noticed it, the data was saved. From this moment, you have a constant, updatable log of competitor activity.

# Chapter 8 Summary

## Main Automation Entities

### *Bots*

1. Suitable for event-driven scenarios with rigid logic.
2. Help automate data collection, communication, etc., based on the "if → then" principle.
3. Work best where the process is predefined.

### *RPA*

1. Used to replace human manual actions in interfaces.
2. Allows automation of legacy IT landscapes without code modifications.
3. Effective in large organizations but requires investment in implementation and maintenance.

### *AI Assistants*

1. Accelerate human work (analysis, information search, help with text, code).
2. Work on user request.
3. Good for improving decision quality.

### *AI Agents*

1. Capable of acting autonomously.
2. Suitable for complex multi-step processes and scenarios with minimal human involvement.
3. Require clear instructions, memory management, and error control.

## **Tool Classes**

### *Bot Builders*

1. A quick and safe way to automate communication and simple scenarios.
2. Best suited for marketing and service processes.

### *Visual Integration Platforms*

1. Allow connecting different services into a single process without code.
2. Optimal for automating typical business chains and event-driven processes.

### *RPA Platforms*

1. Used where there are no integrations or they are too expensive.
2. Provide quick effect in regulatory processes but are sensitive to interface changes.

## **Practical Conclusion**

Automation is a way to systematically reduce manual labor and increase process manageability. The key mistake is expecting a tool to perform a role it's not designed for.

Effective automation starts not with choosing a platform, but with understanding what exactly you are automating: actions, thinking, or decisions.

# Conclusion

With this I conclude the second theoretical-practical part of the book, where I tried to show a wide spectrum of AI tools. In this part I purposely brought together everything we discussed earlier: data, documents, visualization, AI, and automation. So that you understand these are all elements of one system. In real business, no one automates "data cleaning" or "dashboards" just for the sake of it. They automate processes: decision-making, control, communication, reducing manual labor.

And here it was important to dispel the main myth – about "everything turnkey." Over the years, I've seen too many times how beautiful platform promises crash against reality: messy data, undefined processes, blurred responsibility. Therefore, my task in these chapters was not to sell you another tool, but to provide reference points that allow you to soberly assess – what even makes sense to automate, and what doesn't yet.

In this section, we've covered quite a journey – we learned to use AI for typical yet critically important tasks and saw where AI truly saves time.

My task was to convey the importance of learning to think differently, more broadly, in terms of your goal. Not pointwise "how to write a neural network prompt," but "how to build a process with its help": define the audience, ask the right questions, choose sources, understand the degree of information condensation, and formulate conclusions. The same applies to automation in general.

Therefore, another conclusion I want to lead you to: automation is less about technology and more about **manageability**. It starts not with choosing a platform or AI, but with an honest answer to the question: what exactly are we automating – actions, thinking, or decisions – and why. Everything else is secondary.

AI will develop, tools will change, new services for visualization, documents, and agents will appear. But the human role in this system will remain key: to see the meaning, set the boundaries, check the result, and take responsibility. If you understood that – then this part of the book was not written in vain.

# PART 3. Workshop: Case Studies

## Analysis

Until now we have talked about technologies, approaches, and differences between tools. But the main shift that AI brings is not technological but role-related. A person is less often an executor and increasingly **an architect of solutions**. Therefore, it's worth perceiving practice as a change in the human role. Research by McKinsey and BCG has established that companies achieve sustainable benefits from AI not when they implement the "smartest models," but when they start applying them in specific, understandable processes – calculations, reports, planning, and decision-making. It is precisely such tasks that we will analyze next.

This is the final part of the book. In it I have compiled three comprehensive case studies that show how AI takes on routine and preparatory work while a person focuses on meaning and choice. All three cases you will see can be perceived as mental templates.

**The first case:** a calculator for preliminary employee bonus calculations. Perhaps you'll want to replicate this case for your own team. Or perhaps it will become the

basis for a similar but your own solution. The main thing here is the approach: how AI helps quickly assemble a working model and test managerial hypotheses.

In **the second case** we will analyze in detail the creation of an HTML dashboard with AI. That very dashboard that analysts used to spend months on. They manually cleaned data, coordinated visualizations through a long bureaucratic chain, and assembled everything piece by piece. I will show how this process can be accelerated several times and done independently without involving specialists from other departments. At the same time the quality and effectiveness of such a dashboard are not inferior to solutions built in classical BI systems.

**The third case** will be about developing a business plan with AI. We will ask AI to structure our idea, see risks, and justify decisions for investors. A well-made business plan directly affects the trust in the project and the likelihood of its support. In this case, we focus not on what results AI can produce, but on how to formulate requests correctly to make those results useful. The results of this case are the actual work of one of our students from the neural networks course. We separately requested permission to publish their project and included this example in the book exactly as it was done during the course.

I would like to add that this part is not an instruction or a step-by-step manual. I will not hold your hand and show you where to click and what to choose. After all, the goal of this book is to help you understand more deeply how AI works and to reconfigure your thinking for the new reality.

# Chapter 9. Bonus Calculator for Employees

To boost motivation among managers in the team the head decided to create a bonus calculator for them. This panel resembles a dashboard but it's not quite one. A dashboard would be overkill for such a task. However, a visual calculator where managers could see for themselves how they could influence their earnings would be very handy and could have a more impressive effect than a regular presentation. This tool would also be useful for each manager individually. Therefore, it was decided to build such a calculator in HTML format. To create the panel, we used Claude.

## Initial Data

The initial data was an Excel file with two sheets. On the first sheet, managers recorded their sales daily and tracked plan fulfillment. The second sheet contained aggregated monthly data and was maintained by the manager.

**Sheet 1 "Actual Sales"** (1810 rows) – daily manager entries:

- Manager, Full Name
- Date
- Actual Sales
- Units Sold
- Category (Furniture, Books, Electronics, etc.)
- Sales Channel (Distributor, Wholesale, Marketplace, Retail)

**Sheet 2 "Bonuses and Plans"** (60 rows) – monthly aggregates:

- Manager
- Month
- Sales Plan
- Actual Sales
- Bonus
- Base Salary
- Total Returns Amount
- Returned Units

# Data Preparation and Transformation

To build the HTML dashboard the data was merged into a single CSV table (how to do this with AI – I explain in detail in the next Chapter 10). The CSV format was chosen because it's easier to store and load large datasets. Instead of two separate sheets we needed a single table with information for each manager: whether they met the plan, their bonus, any salary changes, the number of returns...

The final table (60 rows × 23 columns) included:

- Manager, Month, Actual\_Sales, Sales Plan, Bonus, Base Salary;
- 10 product categories (expanded into separate columns);
- 4 sales channels: Marketplace, Distributor, Wholesale Trade, Retail Trade.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Sales Manager	Month	Actual Sale	Units Sold	Target Bonus	Base Salary	Returns	As Returned	Units	Category: Clothi	Category: Cosme	Category: Garder	Category: Groceri	Category: Toys
2	Anna Mikhailova	01/2023	5283301	164	5023888	29255	90000	245091	8	9	7	5	10	4
3	Anna Mikhailova	02/2023	2865413	134	3433822	18923	90000	183668	6	7	22	20	15	6
4	Anna Mikhailova	03/2023	4559142	148	4312185	18795	90000	183686	6	2	7	6	11	17
5	Anna Mikhailova	04/2023	4345955	140	4132115	18104	90000	120001	4	23	7	22	5	9
6	Anna Mikhailova	05/2023	4322797	140	4106657	21057	90000	205960	6	14	27	10	31	0
7	Anna Mikhailova	06/2023	2945891	135	3745713	18985	90000	157103	5	15	14	20	11	8
8	Lev Smirnov	01/2023	2100394	142	2002974	16595	90000	101517	8	20	25	8	19	23
9	Lev Smirnov	02/2023	2451918	157	2332172	20645	90000	141804	8	17	36	14	21	19
10	Lev Smirnov	03/2023	2540096	156	2413000	10005	90000	140792	8	7	16	2	13	8
11	Lev Smirnov	04/2023	2283314	159	2174877	20065	90000	87114	5	13	11	17	30	12
12	Lev Smirnov	05/2023	2242220	139	2100117	21129	90000	62332	4	17	19	7	7	10
13	Lev Smirnov	06/2023	2193373	139	2085604	16082	90000	44557	3	1	22	36	6	27
14	Emma Ivanova	01/2023	3045474	166	2833199	17172	75000	172818	8	22	3	2	35	10
15	Emma Ivanova	02/2023	3120725	155	2654689	14730	75000	77521	4	31	11	11	28	12
16	Emma Ivanova	03/2023	3138241	164	2976509	16805	75000	125426	6	14	21	12	11	1
17	Emma Ivanova	04/2023	4086438	163	3882116	15192	75000	101772	5	12	3	24	44	4
18	Emma Ivanova	05/2023	2258331	147	2154942	13859	75000	175015	8	25	10	18	5	26
19	Emma Ivanova	06/2023	3658738	153	3361801	13972	75000	118528	6	0	10	31	13	8
20	Isabella Tikhonova	01/2023	4540724	154	4317688	15919	70000	192652	7	38	16	15	23	12
21	Isabella Tikhonova	02/2023	3465378	148	3290209	12912	70000	194459	7	33	13	28	0	2
22	Isabella Tikhonova	03/2023	5169474	155	4784850	15761	70000	24512	4	15	29	18	13	4
23	Isabella Tikhonova	04/2023	4000324	142	3800300	16005	70000	162632	6	22	8	11	6	9
24	Isabella Tikhonova	05/2023	4317295	170	4152783	16981	70000	182460	7	11	28	37	38	5
25	Isabella Tikhonova	06/2023	3399390	146	3219004	17947	70000	115255	4	1	21	22	13	22

To preserve data confidentiality, the dataset was populated with distorted information while maintaining the structure.

## Creating the Calculator

First, we determined what charts and data we wanted to see on the calculator. We wanted to place cards for plan, actuals, and sales forecast at the top, the "What-if" calculator itself and a bar chart comparing plan vs. actual per manager below. The idea was that by logging into this calculator, managers could view the current month and see how their bonus payout is shaping up. We also decided it would be great if they could see results from the previous month.

To have Claude implement these ideas, we used this prompt:

I want you to build me a dashboard. Format 16:9. Use the latest CDN Plotly.js and include a file upload function so we can load a CSV file. It should generate a monitoring panel based on the uploaded data. Based on my dataset, please help create a dashboard.

Row 1 (top): three KPI cards

1. **\*\*Plan\*\*** - Planned Revenue = Sales Plan × Average Ticket - Sales Plan = sum of the \*Sales Plan\* field - Planned Bonus = (Planned Revenue / Actual Revenue) × Actual Bonus
2. **\*\*Actual\*\*** - Actual Revenue = sum of \*Actual\_Sales\*  
- % Completion = Actual Revenue / Planned Revenue  
- Actual Bonus = sum of the \*Bonus\* field

3. **Forecast** - Forecast Sales = slider value - Forecast Revenue = Forecast Sales × Forecast Average Ticket - Forecast Bonus before Penalty = (Forecast Revenue / Actual Revenue) × Actual Bonus - Penalty = max(0; (Forecast Returns - Allowable Returns)) × (Forecast Average Ticket × 0.1) - Final Bonus = Forecast Bonus before Penalty - Penalty

**Row 2 (bottom): two blocks**

**Left** — chart "Revenue by Manager: Actual vs. Plan" - Actual = sum of *Actual\_Sales* per manager - Plan = sum of *Sales Plan* × Average Ticket

**Right** — "What-if" calculator - Sliders: Forecast Sales, Forecast Average Ticket, Allowable Returns, Forecast Returns - Results displayed: - Forecast Sales - Forecast Revenue - Bonus before Penalty - Penalty - Final Bonus.

As a result, we got the first version of the calculator. It had almost everything we wanted: three cards, the calculator itself, and plan/actual comparison per manager. However, the design and charts required further refinement and fine-tuning. That's exactly what we proceeded to do.



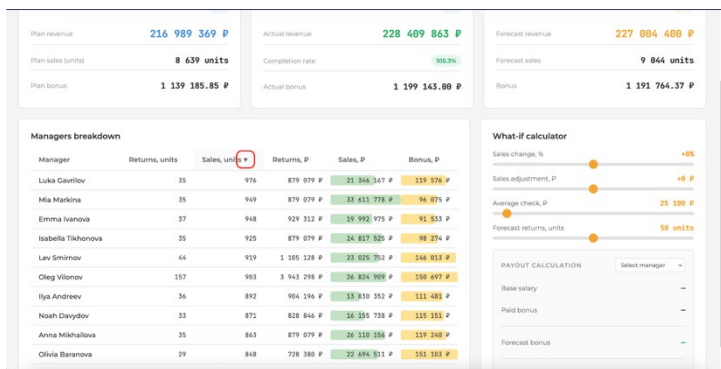
It was decided to use a light theme and refine the calculator. After all, a simple plan vs. actual comparison is often insufficient. To allow each manager to study their situation in detail, it was decided to replace the chart with a table. It was also decided to add a cell on the right in the calculator that would reflect data for only one selected manager. Thus, using the calculator sliders one could see both the overall picture reflected in the Plan, Actual, and Forecast cards for the whole team, and an individual picture for each manager using the employee filter in the right section. The month filter was moved from the block showing forecast results for a specific manager to the top panel.

Initially the neural network assumed there was a certain amount of returns allowable for any manager that didn't

affect the bonus payment. But in reality the company didn't work that way. Any return affected bonuses, and in case of a return the profit the manager earned from selling the product was lost.

Therefore it was decided to remove the slider for allowable returns keeping only the forecast returns. To allow viewing forecasts on larger data in general, it was decided to add sales forecasting in the calculator as a percentage. This allowed not being tied to specific monetary ranges on the sliders.

All this was formulated in a prompt and given as a task to Claude. After these adjustments, the calculator was almost ready. We also added table filtering, with a triangle appearing next to the filtered column name, marking it.

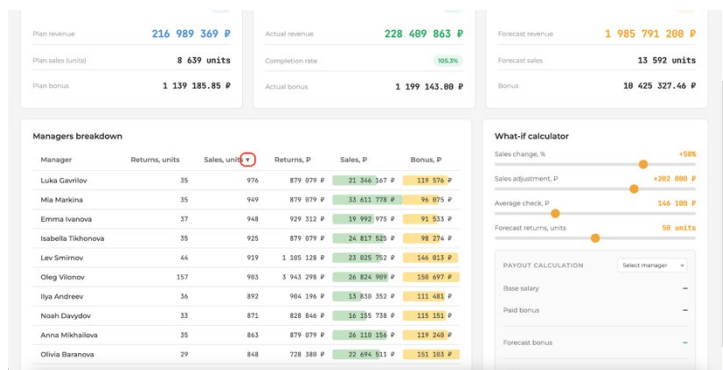


At first glance everything looked correct but the numbers we got seemed unrealistic. This was because we didn't set clear calculation logic. We only described it in general terms, which elements and dependencies existed. Here, we had to discuss all interdependencies with the AI asking it to give an example of how each indicator is calculated. To check the already implemented dashboard, we compared how the calculator computes with our own manual calculations and the results matched.

In the lower left corner, we indicated the calculation logic for the indicators in general terms in case a manager had questions about a specific number.

## Final Version

As a result, we have a quickly and easily created interactive panel.



In the top row we see whether the team's plan is being met overall and the forecast results expected from the team by the end of the month. We see a breakdown for each manager in a table, and we can filter this table.

On the right any manager can study their personal forecasts based on various indicators: sales, average ticket, number of returns.

Such a panel made the manager's presentation more interesting, convincing, and interactive, as each manager can view information personally about themselves, interact with this panel themselves, and model their situation for the current month.

The panel became a tool to return to repeatedly as data changes. All you need to do is upload new information.

Would a presentation have been as convincing and useful?

# Chapter 10. Financial Dashboard

**Task:** Create a full-fledged HTML dashboard that can tell the story of the company "FinExpert Consulting" and serve as an analytical foundation for top management. The dashboard should illustrate the company's results in implementing key projects and suggest where to focus and what to optimize to improve efficiency.

## Initial Data

We have **data from a consulting company**, which consists of several separate exports in Excel file format. There is actual data: "Date", "Projects", "Currency", metrics like "Commercial Expenses", "Cost of Goods Sold", "Revenue". However, the export is done in separate files by year. And the current year (Open Period) has the same table structure but data for months is exported to separate sheets, not into a single array. The analyst's task here is not only to merge this data into one array but also to convert it to a single currency so that, for example, we don't sum euros with dollars or tenge. There is also a table with Targets for the same periods. But there is a peculiarity – the table structure differs from the actuals.

This data needs to be transformed into analytical information on which effective management decisions can be based. Experienced analysts working with, for example, Power BI or even Excel, can handle this task. But this process of "transforming" data from disparate tables into an interactive report useful for business will take considerable time. You need to clean up the data structure, convert values to a single currency, create additional measures without which full-fledged analytics is impossible. Then you need to combine this data in one model via reference tables and relationships. And only then can you start creating the interactive report itself.

Neural networks can already speed up this process. Let's see how Claude solves this task.

# Preparing and Transforming Initial Data

It's best to start the dialogue with the neural network by getting acquainted with the data: attach the source files and ask to describe their structure.

Even with this not-so-complex task, the neural networks performed with varying success. ChatGPT quickly reads the data, correctly recognizes the structure, and even sees upcoming tasks: merging, bringing tables to a uniform format.

Claude analysed the data in Java and showed not only a description of the initial data but also presented a convenient comparative table.

🔍 **Key differences between plan and actual:**

Aspect	Actual files (2022-2024)	Plan file
Structure	Each row = transaction with 3 separate columns for metrics	Each row = one indicator (determined by "Adapter" column)
Currency	Specified for each transaction	All in rubles
Data format	Wide table (3 metrics in columns)	Long table (metric in rows)

Would you like me to create a summary analysis or visualization based on this data?

📄 👍 🗑️ Retry ▾

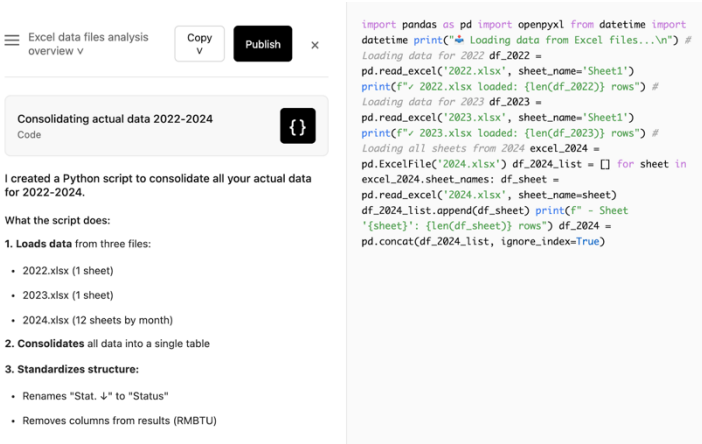
Then we give a more specific task:

Create a single flat table of actual indicators from the 2022, 2023, and 2024 files. Merge all data into one dataset, including data from different sheets for 2024.

The final table should have columns: Date, Project, Item, Currency, Commercial Expenses, Cost of Goods Sold, Revenue.

Save the result in an Excel file.

The neural network didn't output a ready-to-download file on the first try, but only offered ready-made Python code to solve the task of merging files into one array.



Excel data files analysis overview v

Copy v Publish x

Consolidating actual data 2022-2024  
Code

I created a Python script to consolidate all your actual data for 2022-2024.

What the script does:

1. **Loads data** from three files:
  - 2022.xlsx (1 sheet)
  - 2023.xlsx (1 sheet)
  - 2024.xlsx (12 sheets by month)
2. **Consolidates** all data into a single table
3. **Standardizes structure**:
  - Renames "Stat. ↓" to "Status"
  - Removes columns from results (RMBTU)

```
import pandas as pd
import openpyxl
from datetime import datetime
print("Loading data from Excel files...\n") #
Loading data for 2022 df_2022 =
pd.read_excel('2022.xlsx', sheet_name='Sheet1')
print(f"2022.xlsx loaded: {len(df_2022)} rows") #
Loading data for 2023 df_2023 =
pd.read_excel('2023.xlsx', sheet_name='Sheet1')
print(f"2023.xlsx loaded: {len(df_2023)} rows") #
Loading all sheets from 2024 excel_2024 =
pd.ExcelFile('2024.xlsx') df_2024_list = [] for sheet in
excel_2024.sheet_names: df_sheet =
pd.read_excel('2024.xlsx', sheet_name=sheet)
df_2024_list.append(df_sheet) print(f"Sheet
'{sheet}': {len(df_sheet)} rows") df_2024 =
pd.concat(df_2024_list, ignore_index=True)
```

After the next request, we did get a fully ready file with the option to download to our PC.

Now let's proceed to convert currency values. For this we'll use another dataset "Exchange Rates". It's something like a reference where, for a specific date, different currencies, the exchange rate for conversion to dollars, and a multiplier indicator are recorded. We also describe in sufficient detail for the neural network what we want to get in the end:

Use the "Exchange Rates" file to convert all indicators to dollars.

Add three new columns to the actuals table: SG&A USD actual, Cost of Goods USD actual, Proceeds USD actual.

Conversion logic:

If Currency = "USD", no conversion needed, take the original value.

If Currency ≠ "USD", then conversion = Value \* Rate / Multiplier.

Take rates by date and currency from the file «Exchange Rates.xlsx».

Delete the old columns SG&A, Cost or Proceeds.

Save the result in an Excel file "Actuals in USD".

You can continue working with the data in the Claude environment without exporting the file and without using a script.

We move on to Targets. Now the task is to make the table structure the same as for actuals.

Convert the file "Targets 2023-2025" to the same format as the actuals.

In the source data, there is a column Category, indicating whether the amount relates to SG&A, Cost or Proceeds.

Unpivot this column to get a table with columns:

Date, Project, Item, SG&A, Cost or Proceeds.

Important: Replace empty values in the Item column with an empty string (not "No Item").

Ensure Proceeds from the plan indeed ends up in the final table.

Save the result in an Excel file "Target".

As mentioned above you can work with files within the neural network environment itself, previewing transformation results, you can download Excel to your device or you can use the scripts it provides (how to run them, what programs to use – you can also ask the neural network about that). To check the result of the script from Claude let's run its code in a special environment, first asking it to compile a single code that includes everything from merging actuals by years to merging targets and actuals into one document.

Claude, after a couple of requests, produced working code with saving all necessary data.

```
print("\n--- YEARLY PLAN (USD) ---")
print(plan_by_year.to_string())

print("\n" + "=" * 80)
print("PROCESS COMPLETED SUCCESSFULLY")
print("\n" + "=" * 80)

print("\nGenerated files:")
print("  1. Actual_USD.xlsx - actual data in USD")
print("  2. Plan.xlsx - planned data")
print("  3. Actual_vs_Plan_2022_2024.xlsx - consolidated file with two sheets")

print("\nData structures:")
print("  All files contain the following columns:")
print("  - Date")
print("  - Project")
print("  - Category")
print("  - Commercial Expenses")
print("  - Cost of Goods Sold")
print("  - Revenue")
```

=====

**PART 3: SAVING RESULTS**

=====

STEP 3.1: Saving intermediate files

- ✓ File saved: Actual\_USD.xlsx
- ✓ File saved: Plan.xlsx

STEP 3.2: Creating the final file

- ✓ Sheet "Actual" created (5288 rows)
- ✓ Sheet "Plan" created (1712 rows)

Final file saved: Actual vs Plan 2022-2024.xlsx

Thus, the task of preparing and transforming the initial data was solved. You can work with the neural network in different ways: get ready-made Excel files immediately, use the provided scripts, or analyze data directly in the model's interface. You can choose a convenient way of working depending on the task and user experience.

## Creating Measures

Even though building reports with neural networks doesn't require separately creating and outputting additional calculations we want to include in the report, we still included this stage, adhering to the same logic as working with traditional BI tools. This allows seeing how an analyst's routine tasks can be automated to get results much faster. For example, it's important for us to see expense indicators, profit, and percentage of fulfillment.

We go to the neural network with a request where we don't just "demand" to output these indicators, but also explain the calculation logic, i.e., what exactly we mean by this metric and which values from the source data we use.

I have an Excel file with merged and transformed data, containing two sheets: "Actual" and "Target". On both sheets data is in a unified format: columns Date, Project, Item, SG&A, Cost or Proceeds.

I need:

1. Group data by year. Calculate for each year:
  - Actual Expenses (sum of Actual Costs + Actual SG&A).
  - Target Expenses (sum of Target Costs + Target SG&A).

- $\text{Actual Profit} = \text{Actual Proceeds} - \text{Actual Costs} - \text{Actual SG\&A.}$
  - $\text{Target Profit} = \text{Target Proceeds} - \text{Target Costs} - \text{Target SG\&A.}$
  - $\text{Revenue Fulfillment \%} = \text{Actual} / \text{Target.}$
  - $\text{Revenue Deviation} = \text{Actual} - \text{Target.}$
  - $\text{Revenue \% Difference} = (\text{Actual} - \text{Target}) / \text{Target.}$
  - Similarly, calculate Fulfillment %, Deviation, and % Difference for SG&A, Costs, Profit, Expenses.
2. Add last year's values for actuals:
    - Last Year's Actual Revenue,
    - Last Year's Expenses,
    - Last Year's Costs,
    - Last Year's SG&A,
    - Last Year's Profit.
  3. Output all results as a table directly in the chat, rounded to thousands and percentages in %-format.
  4. Briefly explain the calculation logic for each value after the table.

We get the calculations right here in the chat, but if necessary, we can also get them in Excel. At this stage it's still better to "play it safe" and double-check the results, sometimes using traditional tools, sometimes engaging your own critical thinking.

The check showed that both the calculations in the ready file and the results obtained using the Python script completely matched those previously calculated manually using traditional tools.

```

=====
STEP 1: REVENUE DATA AGGREGATION
=====
✓ Aggregation completed

STEP 2: CALCULATION OF PRODUCTION METRICS
=====
✓ All metrics calculated

=====
FINAL TABLE (all amounts in USD thousands, rounded)
=====

1. BASE METRICS (USD thousands)
=====
Year  Revenue Actual  Revenue Plan  COGS Actual  COGS Plan  Commercial Expenses Actual  Commercial Expenses Plan
2022  51,050             45,411        9,233        6,935        318                    351
2023  42,820             32,140        6,572        4,694        308                    377
2024  54,636             65,118        8,123        8,797        764                    554

=====

2. EXPENSES AND PROFIT (USD thousands)
=====
Year  Expenses Actual  Expenses Plan  Profit Actual  Profit Plan
2022  9,543            7,287         42,187        38,144
2023  6,880            4,471        35,940        27,669
2024  8,887            9,352        45,749        55,766

=====

3. PLAN EXECUTION (%)
=====
Year  Revenue Execution %  COGS Execution %  Commercial Expenses Execution %  Expenses Execution %  Profit Execution %
2022  113.7%              133.1%           88.8%                          131.8%              118.4%

```

## Creating the Dashboard

We move on to the final stage – creating the interactive dashboard.

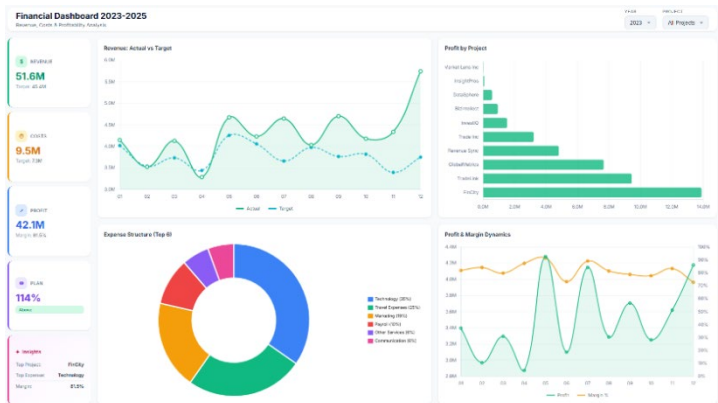
The approach is the same as in Chapter 7 – from developing the idea of the future dashboard to the final version with interactive filters.

We continue working in the same chat where we had the dialogue with the neural network; it's important that the neural network retains the context.

Creating the dashboard itself will also require an iterative approach – you can't create a dashboard with one "magic" prompt.

**First step:** A request to create a dashboard based on the available data (structured and ready).

Claude provides a decent version of the dashboard in terms of meaning and design even without a clear requirements definition. And the report itself even becomes interactive with working filters.



But in this case study we'll take a different path. We already have our own idea of what the final report will look like. Therefore, we describe in detail what needs to be placed on each level of the dashboard, what calculation logic to use.

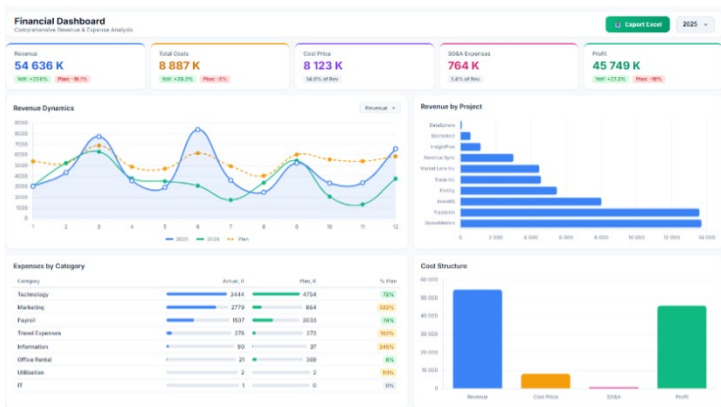
I have an Excel file with two sheets: "Actual" and "Target". On both sheets, data is in the format: Date, Project, Item, SG&A, Cost or Proceeds.. Build a dashboard with the following indicators:

- KPI cards: Revenue, Expenses (Costs + SG&A), Cost of Goods Sold, SG&A, Profit = Revenue – Costs – SG&A. For each metric calculate % deviation to last year and to target.
- Line chart "Revenue Dynamics" by months: actual current year, last year, and target (dashed line). The ability to switch the metric on this chart via a filter – "SG&A Dynamics",

"Expenses Dynamics", "Cost of Goods Sold Dynamics", "Profit Dynamics".

- Horizontal bar chart: "Revenue by Projects" which will switch when using the slicer next to the line chart – to "SG&A by Projects", "Expenses by Projects", "Cost of Goods Sold by Projects", "Profit by Projects".
- Table "Expenses by Category", first column Actual Items, second - Plan (Target) Item, third % Plan Fulfillment. In the second and third columns, use conditional formatting with micro histograms; in the last one, font color green if 100% and above, red if below 100%.
- Waterfall chart "Cost Structure" (Revenue → Cost of Goods Sold → SG&A → Profit).
- Make the dashboard 16:9, focusing on conciseness and clear labels.

The best approach is to form an HTML dashboard with data loading capability from the very first requests. This will allow building the dashboard, for example, not only on test data but also later pulling in current data in the local HTML file, i.e., essentially updating data in the report.



**Second step:** We work on the report's design. With neural networks, you can generate many design options for the report. But, for instance, it's important for us that it follows the corporate style. So we can share with the neural network as a sample: style guides, presentations, the company website – i.e., everything designed in the corresponding style.

We get a version of the designer-style report:



**Third step:** We refine individual visuals (cards, charts). Similarly to other case studies, we specify in the prompts how each visual should look. For example, charts should have data labels enabled, and scales with approximate values, which currently only create extra informational noise, should be disabled. We configure a unified style for report fonts, "agree" on what the color coding for indicators will be.



At the final stage of dashboard creation, we add a scenario modeling block "What if...". Its goal is to give the user the ability to quickly assess the impact of changes in key indicators on the financial result. Under the main KPI cards, a separate interactive block with cards is formed, where for revenue, cost of goods sold, and SG&A, you can set a deviation percentage within an allowable range. All values are recalculated instantly and the resulting values – total expenses and profit – are updated automatically based on the defined calculation logic. This approach allows, without complex formulas and external tools, to quickly test various management scenarios directly within the report.



As a result, we get not just a beautiful report, but a convenient working tool. The dashboard has all key indicators, filters, and clear charts, as well as a "What if..." block, which allows quickly seeing how changes in revenue or expenses will affect profit. The neural network helps step by step to assemble such a dashboard – from the general idea to fine-tuning calculations and appearance. As a result, one report can be used for both regular analysis and management decision-making, without switching between different files and tools.

## Chapter 11. Business Plan

Imagine you have a brilliant business idea and you need to present it to investors tomorrow. Besides the idea you have passion and a lot of enthusiasm, but you still lack a coherent financial model, clear answers to questions about payback period, customer acquisition cost, or product lifespan on the market. The task of any businessperson is to see opportunities, not just limitations. But how to look at a problem from a different angle when your knowledge and experience limit your perception?

Most start-ups and entrepreneurs face exactly this. Even if you're not a finance expert and don't plan to delve deep into economic formulas, you'll have to speak the language of numbers and metrics at the meeting. This is where AI comes to the rescue: it can quickly explain what LTV and CAC are, help compile an expense table, and highlight weaknesses in your model. But (and this is the key point) it will not replace your understanding of the logic and essence of the indicators.

In this chapter, we will break down the basic approaches to business planning and financial modelling as if we are beginners. Why is this needed? So that AI becomes your amplifier, not a source of errors. You will learn to

ask it the right questions, check the answers received, and understand when it's worth double-checking calculations manually.

This case study will be half theoretical: we will lay the foundation, break down key concepts, and understand why they are needed at all. At the same time, we'll apply this to practice - we'll see what results AI gave during our students' training.

### **Minimalist Approach: Finance Without Frills**

The presentation on this topic should appropriately be done in a minimalist style – as befits finance professionals. What's important here is not visual appeal, but logic and calculations. We are not trying to amaze with a slide; we are proving an idea with numbers.

# Classical Business Planning Model

## APPROACHES TO PLANNING



### TRADITIONAL (OFFLINE)

FOCUSED ON CAPEX (CAPITAL EXPENDITURES), OPEX (OPERATING EXPENSES), AND CASH FLOW MANAGEMENT.



### DIGITAL

BUILT AROUND UNIT ECONOMICS: CALCULATING THE EFFICIENCY OF A SINGLE CUSTOMER AND SCALING ACQUISITION.

## Traditional Business Planning

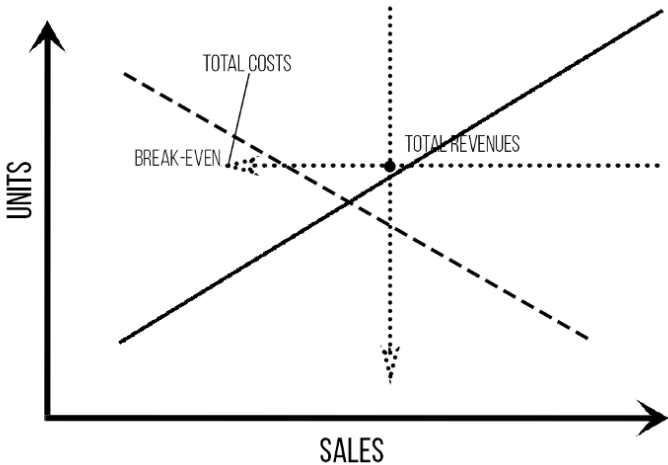
Let's start with the traditional approach. This model is familiar to many who studied finance in university or encountered preparing documents for an investor:

1. Assessment of Capital Expenditures (CAPEX): Equipment purchase, renovation, licenses.
2. Operating Expenses (OPEX): Rent, salaries, marketing.
3. Revenue Forecast: Based on price and sales volume.
4. Performance Indicators: Payback period, ROI (Return on Investment), break-even point.

## KEY FORMULAS AND RULES

<b>ROI</b>	=	$\frac{\text{NET PROFIT}}{\text{INVESTMENTS}} \times 100\%$	<b>BENCHMARK:</b> ROI > 20% ANNUALLY
<b>PAYBACK PERIOD</b>	=	$\frac{\text{INVESTMENTS}}{\text{AVERAGE ANNUAL NET PROFIT}}$	<b>EXPECTED PERIOD:</b> 2-3 YEARS
<b>EBITDA MARGIN</b>	=	$\frac{\text{EBITDA}}{\text{REVENUE}} \times 100\%$	<b>TARGET:</b> 15-30%
<b>BREAKEVEN POINT</b>	=	$\frac{\text{FIXED COSTS}}{1 - \text{SHARE OF VARIABLE COSTS IN REVENUE}}$	<b>ACHIEVED IN:</b> 12-18 MONTHS

*ROI above 20% is a good benchmark for small business. Lower is already less interesting for investors.*



If you plan to open a cafe, a studio, or another local business, it's important for you to calculate: in how many months will the investment be recouped. For example, promises of a 6-month payback need to be supported by calculations – not faith.

Try calculating the payback period yourself providing the initial data to AI (for basic calculations, you can use any neural network you prefer). Use realistic estimates of the average check in that particular business you'd like to practice on.

I've talked about the traditional approach to business planning; now let's break down the essence of digital planning based on unit economics.

## Digital Business Planning (Unit Economics and Digital Products)

In IT and online educational projects the financial model is structured differently. Here, unit economics comes into play: metrics calculated per customer.

### **Key Indicators:**

- CAC (Customer Acquisition Cost): How much it costs to acquire one customer.
- LTV (Lifetime Value): How much money a customer will bring over their lifetime.

- Churn Rate (Attrition): What percentage of users leave.
- Payback Period: After how many months a customer pays back.

Formulas for calculating the listed indicators:

#### CUSTOMER ACQUISITION COST

$CAC = \text{MARKETING BUDGET} / \text{NUMBER OF ACQUIRED CUSTOMERS}$

#### LIFE TIME VALUE

$LTV = ACV \times \text{NUMBER OF PURCHASES PER CUSTOMER LIFETIME} \times \text{GROSS MARGIN}$

#### FOR SUBSCRIPTION MODELS

$LTV = \text{AVERAGE MONTHLY REVENUE} / \text{CHURN RATE}$

#### ESTIMATING PAYBACK PERIOD

$PAYBACK = CAC / \text{AVERAGE MONTHLY PROFIT PER CUSTOMER}$

#### ANALYSIS OF

RETENTION AND CHURN

**Rule:** LTV should be 3 times greater than CAC – then the model is considered sustainable.

## KEY RULES

LTV > 3 × CAC — MODEL SUSTAINABILITY

ACQUIRE A CUSTOMER FOR 50K AND GENERATE 120K — THAT'S A FAILURE

PAYBACK PERIOD < 6 MONTHS

FAST PAYBACK OF CUSTOMER ACQUISITION INVESTMENTS

1-MONTH RETENTION > 40%

ACCEPTABLE LEVEL FOR GROWTH

In practice it's different. One EdTech company failure case: they scaled advertising, spent hundreds of thousands on acquisition, but customers didn't pay back. It seemed like everything was growing, but in the end, there was a cash gap — they simply didn't have enough money to sustain growth.

### **Retention as a Critical Factor**

In digital business, it's important not only to attract but also to retain. If a customer left after a month, you didn't have time to recoup them. This is especially critical in a subscription model. If churn is high — the project won't survive even with active advertising.

It is considered normal when a customer starts paying in the first month and continues to use the service.

Therefore, retention becomes almost more important than acquisition.

### Comparison of Approaches

<b>Parameter</b>	<b>Offline Business</b>	<b>Digital Product</b>
<b>Main Task</b>	Return on Investment	Marketing Efficiency
<b>Key Metric</b>	ROI, Payback Period	LTV/CAC, Retention
<b>Planning Horizon</b>	2-3 years	6-18 months
<b>Scaling</b>	Through Physical Expansion	Through Traffic Growth

When preparing a business plan, it's important to understand the difference between approaches to business valuation and skillfully combine the necessary estimates and metrics from both paths.

# Creating a Business Plan

The director tasked us with preparing an investment business plan for a project worth about \$1 million.

Our goal is to convincingly show that we are ready to scale the production of vertical solar panels and effectively utilize the investments.

A prompt was written for the neural network to create a business plan and presentation, specifying the main initial data about the project:

Intro about our project:

We are the company SOLARWALL producing flat vertical panels for facade and balcony mounting.

The project is dedicated to launching the production of vertical solar panels for installation on balconies and building facades. In recent years the solar energy market has shown stable growth of about 20% per year, and one of the current trends has been the development of solutions for urban environments, where space for traditional horizontal panels is limited. Vertical panels allow individuals and small businesses to install solar modules even in small areas, reducing electricity costs and promoting the greening of urban environments.

Features of our panels:

- Panel size: ~1200×500 mm
- Power per panel: 150 W

- Frame – aluminum, linear
- Assembly – in a small workshop (garage format)
- Standard fasteners for straight walls
- Assembly time per panel: 10-12 hours

We have been on the market for 1.5 years and have over 200 clients from different segments: Private Individuals, Corporations, Homeowners Associations, Management Companies, Small Businesses.

What needs to be done:

1. Compile a business plan:

- Description of the market, target segments, competitive environment.
- Production model before and after investments.
- Financial model: cost, revenue, profit, break-even points.
- Plan for utilizing investments.

2. Create a stylish presentation for the investor:

Briefly, clearly, and convincingly present the essence of the project, its prospects, and how exactly the investments will be utilized.

This prompt is quite brief but comprehensive. It's quite sufficient to "inspire" the neural network if you use an up-to-date model – for example, GPT-5.2 at the time of writing this chapter. A simple piece of advice here: choose the model with the highest level of reasoning

(**Model reasoning** is the process where an AI model "reasons" to arrive at an answer to the user's question).

Here, I'll present its table of contents:

PRODUCTION OF VERTICAL SOLAR PANELS FOR FACADES AND BALCONIES .....	1
1. Project Summary .....	1
2. Market Opportunity .....	2
3. Product .....	2
4. Production .....	3
5. Demand Validation .....	3
6. Marketing and Sales Funnel .....	4
7. Financial Model (12 Months) .....	4
8. Loan Repayment and Collateral .....	5
9. Scenario Planning and Adaptation .....	5
10. Team .....	6
Appendices .....	6
APPENDIX 1. TECHNICAL PRODUCT DESCRIPTION .....	7
APPENDIX 2. FINANCIAL MODEL: THREE GROWTH SCENARIOS .....	9
Base Scenario (Planned) .....	9
Pessimistic Scenario (Low Demand) .....	9
Optimistic Scenario (High Demand) .....	10
APPENDIX 3. PRODUCTION AND PROCUREMENT PLAN .....	11
APPENDIX 4. MARKETING STRATEGY .....	13
APPENDIX 5. LOAN REPAYMENT SCHEDULE .....	14
APPENDIX 6. LOAN REPAYMENT SECURITY .....	17
APPENDIX 7. MARKET AND COMPETITIVE ANALYSIS .....	19
APPENDIX 8. RISKS AND MITIGATION PLAN .....	20

The report ended up being 21 pages long. Of course, several iterations are needed to ensure that GPT added all necessary information and the report looks truly

professional (for example, emojis definitely have no place in business documents), the structure is logical, the project goal is clearly formulated, and conclusions are drawn based on the business plan steps and action roadmap.

Describe exactly what financial models you want to see in the business plan. Typically, it makes sense to include several development scenarios: optimistic, realistic, and pessimistic. This allows you to assess in advance how the project's economics will change under different conditions – growth or decline in revenue, changes in costs, payback periods, and key metrics. This approach helps not only to better understand business resilience but also to prepare for questions from investors and partners. Here is a fragment of our business plan:

### Base Scenario (Planned)

- Moderate sales growth
- The entire batch is sold within 6 months
- Break-even is reached within 2-3 months

Month	Units Sold	Revenue (\$)	Cost of Goods Sold (\$)	Gross Profit (\$)	Operating Expenses (\$)	Operating Profit (\$)
1	30	9,000	3,900	5,100	7,000	-1,900
2	45	13,500	5,850	7,650	7,000	+650
3	60	18,000	7,800	10,200	7,000	+3,200
4	60	18,000	7,800	10,200	7,000	+3,200
5	55	16,500	7,150	9,350	7,000	+2,350
6	50	15,000	6,500	8,500	7,000	+1,500
<b>TOTAL</b>	300	90,000	39,000	51,000	42,000	Net operating profit: +9,000

### Pessimistic Scenario (Low Demand)

- Slow start and weak market response
- The batch is sold over 9 months
- Cost optimization and expansion into additional sales channels are required

Month	Units Sold	Revenue (\$)	Cost of Goods Sold (\$)	Gross Profit (\$)	Operating Expenses (\$)	Operating Profit (\$)
1	20	6,000	2,600	3,400	7,000	-3,600
2	25	7,500	3,250	4,250	7,000	-2,750
3	30	9,000	3,900	5,100	7,000	-1,900
4	35	10,500	4,550	5,950	7,000	-1,050
5	35	10,500	4,550	5,950	7,000	-1,050
6	40	12,000	5,200	6,800	7,000	-200
7	40	12,000	5,200	6,800	7,000	-200
8	40	12,000	5,200	6,800	7,000	-200
9	35	10,500	4,550	5,950	7,000	-1,050

And, of course, we can't do without marketing strategy. Like additional content of the business plan (for example, financial model scenarios, collateral, and loan repayment schedule, etc.), it is better configured gradually, through iterations, to avoid overwhelming the neural network with too much information. AIs are generally prone to getting confused in their own statements if you dump too many inputs on them at once.

I usually recommend describing the entire process as a whole but strictly limiting the transition to the next step with commands. Remember how we did this in the practice with the board of directors?

Using the same principle through "rounds," it is convenient to assemble a business plan or any other lengthy and important documentation. This approach helps keep logic under control and get a clearer result.

#### Appendix 4. Marketing Strategy

##### Target Audience

Segment	Description	Purchase Motivation
<b>Urban Apartment Residents</b>	Apartment owners with balconies or loggias	Electricity cost savings, sustainability, energy independence
<b>Small Businesses (cafés, workshops)</b>	Use panels on facades, storefronts, and terraces	Cost reduction, "green business" brand image
<b>Suburban Homes, Garden Associations</b>	Install panels vertically (fences, walls, facades)	Easy installation, no roof required

##### Product Positioning

"A solar panel you can install yourself. No permits required. No electrical work."

Key Advantages:

- Easy installation — no specialists required
- Aesthetic and compact, adapted to urban architecture
- Flexible installation options: facade, balcony, fence
- Local manufacturing — lower cost compared to imported alternatives

##### Marketing Channels (Initial Phase)

Channel	Format and Objectives	Budget (First 3 Months)

There is another fundamentally important point – to do a spot check, i.e., selective verification of calculations. Yes, AIs are becoming more accurate in calculations,

but they still need to be double-checked. You probably wouldn't want to embarrass yourself in front of investors by presenting a final ROI value that turned out to be significantly lower or higher than the real one. Investors don't like inaccuracies.

For example, check the calculation from similar parts of the compiled business plan (table fragment).

 **Monthly Payment Structure**

Month	Payment, \$	Interest, \$	Principal Repayment, \$	Outstanding Balance, \$
1	758	758	0	65,000
2	758	758	0	65,000
3	758	758	0	65,000
4	758	758	0	65,000
5	758	758	0	65,000
6	758	758	0	65,000
7	4,080	760	3,320	61,680
8	4,080	718	3,362	58,318

Always cross-check numbers with source data and formulas used by GPT to calculate certain indicators. In the initial prompt you can even ask to provide the formula used in the answer to ensure the variables are correct.

# Creating a Presentation

It might seem that with AI you can solve any task in a couple of clicks. But reality is like that: quality is not just about technology, but also brainwork. Now we will test Gamma.app to create a presentation based on our business plan.

Was it quick? No. But the result was worth the effort. We share our experience – step by step, briefly, so you can replicate it and get the same great result.

## Step 1. Loading Data into AI

The first step seems simple: take the obtained business plan, upload it to GPT, and write *"Highlight the most important key points from the document for creating an investor presentation."*

But this is where the challenge lies: simply writing a request like *"Make it look nice"* will bring disappointment – fluff and some generalities.

**The correct way:** The quality of the answer is 90% dependent on the quality of your prompt. We used the G-C-S-F-S formula (Goal-Context-Source-Format-Style).

## THE FORMULA FOR THE PERFECT PROMPT



WRITE 7–10 KEY TAKEAWAYS THAT EMERGED DURING DISCUSSIONS WITH THE CLIENT ABOUT THEIR FUTURE BRAND.

THESE SHOULD HELP ME PREPARE FOR A MEETING WITH THE CLIENT TO DISCUSS THEIR ADVERTISING CAMPAIGN.

FOCUS ON THE ATTACHED EMAIL CORRESPONDENCE FILES.

PRESENT THE TAKEAWAYS IN A TABLE WITH BRIEF EXPLANATIONS.

USE SIMPLE LANGUAGE SO I CAN QUICKLY GET UP TO SPEED ON THE TOPIC.

As a result:

### 1) Project Overview and Objective

- **Project:** Mass production of **vertical solar panels** designed for installation in **limited urban spaces** — **balconies and facades** (residential buildings and small commercial properties).
- **Objective:** To provide consumers with an affordable solution to **reduce dependence on centralized energy sources** and enable more **sustainable energy consumption**.
- **Product Format:** A ready-to-use kit designed for **self-installation**, without the need to involve professional installers.

### 2) Market Opportunity and Problem

- **Market Pain Point:** In urban environments, there are virtually no accessible solutions, as traditional solar panels are primarily designed for rooftops. Many existing options require permits, complex installation, and dedicated engineering infrastructure.
- **Trends:** Growth of the solar energy market (**+20% CAGR**), rising electricity tariffs, increasing interest in **energy efficiency / sustainability / autonomy**, as well as urban programs focused on **decarbonization** and “**green facades**.”
- **Target Segments:** Apartment owners with balconies, small commercial properties (cafés, workshops, retail spaces).



## Step 2. Analyze and Clarify

AI is not a god and its answers require critical assessment. But isn't it the same with people? We don't take colleagues' and experts' words at face value – we check, analyze, clarify. With AI it's all the same. It doesn't replace thinking; it offers versions that require thoughtful verification.

We went through each point:

- Removed unnecessary points.
- Asked AI to expand suitable points, adding more details.

As a result, we got the final text for the presentation.

**Advice:** At this and every subsequent stage it's important to critically evaluate the conclusions. You are the expert; AI is your assistant.

## Step 3. A Pinch of Creativity (Optional)

The working text is now there, but it lacks engagement. How to fix it?

Again, we turn to AI for ideas: \*Suggest how to make our presentation more interesting, more memorable, and increase the likelihood of a positive response from investors?

## We got this answer:

Short and investor-focused — what strengthens the pitch and increases the chance of a “yes”:

**1. Start with context, not the product.**

Slide 1: *“Urban solar energy today is either too complex — or simply useless.”*

The investor immediately understands **why the market is broken**.

**2. Craft one memorable sentence.**

Example format: *“A solar panel that finally works for the city.”*

This is the anchor the investor will walk away with.

**3. Show traction before financials.**

Real applications, pre-orders, LOIs — **before the tables**.

Removes the core fear: *“Who actually needs this?”*

**4. Make the financial model “human.”**

Not spreadsheets, but three facts:

how much one panel costs, how much we earn per unit, and from which month the project stops needing cash.

**5. Clearly define what the investor is actually buying.**

Not “panels”, but a production model.



The idea is simple, but we think this option can add liveliness to our presentation.

## Step 4. Assemble the Document Structure and Upload to Gamma App

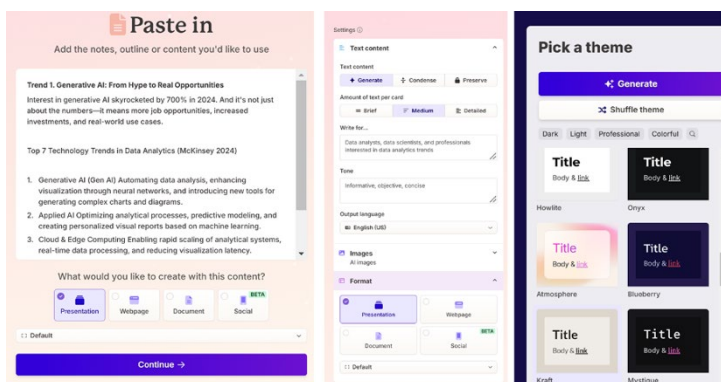
We ask AI to create a final document based on a basic structure:

## PART 1. BASIC STRUCTURE OF AN INVESTMENT PRESENTATION

1. Cover / Title
2. Market Problem
3. Why Now (Market Timing)
4. Solution / Product
5. Customer Value (Value Proposition)
6. Market and Target Audience
7. Traction / Proof of Demand
8. Business Model
9. Unit Economics
10. Production and Scalability
11. Go-to-Market Strategy
12. Competitive Landscape
13. Financial Model (High-Level)
14. Growth Scenarios
15. Investment Ask and Use of Funds
16. Team

Overall the file is ready. I recommend manually going through the text once more and refining it if necessary.

Gamma App is a tool that automatically turns text into stylish presentations or landing pages.



## How is Gamma different from basic PowerPoint?

- Speed. Gamma converts pasted text into a ready presentation in a couple of minutes.
- Convenience. Intuitive interface without unnecessary tools; it's a simple constructor for your slides.
- AI Integration. Gamma helps structure material even at the creation stage. It can generate new text for you, rewrite old text, create images based on your request.
- Online Access. Your presentations will be available online via a link and immediately adapted for both computer screen size and mobile phones.

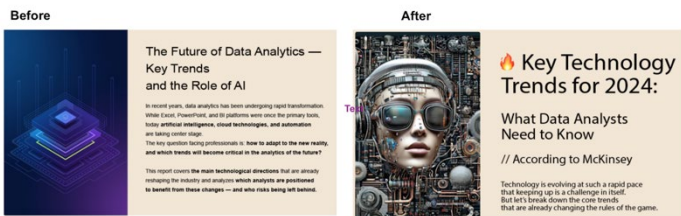
The service is free, but for auto-generating slides based on text or generating AI text from scratch you need "credits." Usually bonus credits are more than enough for initial projects.

### **How to work in it:**

Upload the prepared prompts, choose a design style, and see what you get.

At this step, you may face an obvious fact: automatic generation is just a starting point. It's hard to call the first

version perfect – the slide structure turned out okay, but the text on them needed refinement, and the visual style needed additional polishing.



## What to consider:

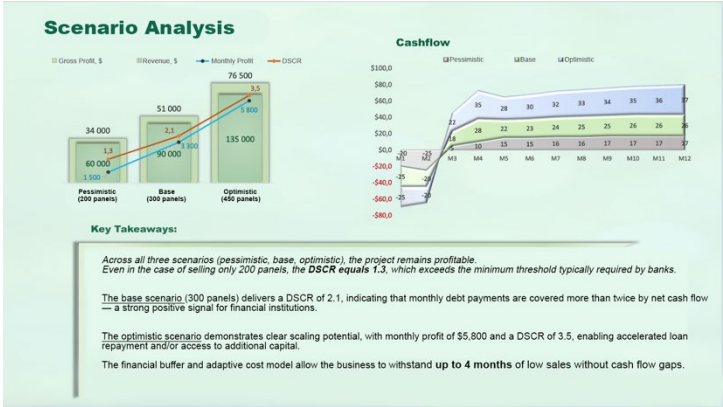
Gamma is great at basic layout but the design requires refinement to look professional. Gamma doesn't replace a designer; it creates a base that you improve.

## Step 5. Polish the Result

Working with AI is about iterations. Gamma generated a draft but it needed refinement. We revised the structure, highlighting key accents, corrected the text to make formulations more precise, and replaced some images, choosing more suitable ones.

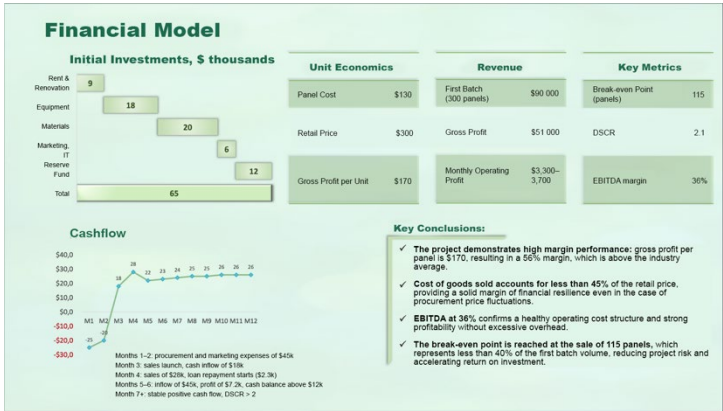
In the end, we have a good result, but this process once again confirmed: AI is a tool, not a magic button. It speeds up work, but quality depends on you.

Based on our business plan case study let's look at the final slides generated based on the document we uploaded to the neural network. *Note:* the amount of manual work here is minimal.



*Work by student Olga Manzhava*

And now an honest question – how much time would you spend on 7–10 similar slides for investors if you made them from scratch manually? In Gamma it would take about 15 minutes. If you're a perfectionist – maximum half an hour.



Of course, as a perfectionist, I would refine the design, typography, and remove visual noise. And now a small independent task: look at these slides and think about what exactly you would change or improve.

# Conclusions

This experiment clearly showed that using AI in analytical work is not just pressing one button, but a skilled interaction between algorithm and expert.

**The main conclusion:** Automatic tools significantly accelerate the process, but the quality of the result still depends on critical thinking and human refinement. AI itself doesn't understand context, and without your intervention, the result will be either overly generalized or simply not fit for purpose.

What is especially important when working with AI:

- Clear task formulation. The more detailed the prompt, the better the result.
- Critical assessment of answers. Always check and clarify the proposed conclusions.
- Iterative approach. AI rarely gives a perfect result on the first try; improvement requires several iterations: clarifying prompts, filtering data, making adjustments.
- Final polishing. Even if AI did excellently, the final edit still remains with the human.

Using AI in analytics and business becomes a powerful advantage, but only for those who know how to work with it consciously.

And the most important thing: there is currently no single corporate standard. Some do everything in ChatGPT, others in Notion, some create visuals in Midjourney, and assemble the final report in Canva. And that's okay. Because AI is like Microsoft Office: you have Word, Excel, PowerPoint – and no one expects one program to do everything perfectly. It's exactly the same with neural networks: each tool solves its own task, and the task of a manager or you as an employee – is to decide whom to assign what or what and where to do.

# Conclusion. How Businesses and Leaders Can Avoid the AI Trap

While working on this book I had the feeling that "the world has radically changed" or that old principles no longer work. New tools, new models, new opportunities have appeared. But, in reality, if you look deeper the basic things remain the same. We still work with people, build processes, make decisions, and are responsible for results, and AI has been added to this system as a powerful amplifier.

To sum up the entire journey we've taken together, we understand that neural networks are still far from being used by everyone and that's normal. Businesses have many objections, doubts, and fears, and I deliberately did not try to dismiss them. In this book we stood alongside the reader, not preached from above. Yes, there are risks. There are illusions and hype. There are situations where AI is used for the wrong purpose, leading to problems. For example, when "neuro-employees" are introduced into a sales department without control and they start promising clients unnecessary things, responding inappropriately, or destroying trust. In such cases not only sales but overall profit declines. But there are also opposite examples – companies that carefully integrate AI into processes, understanding its role, limitations, and area of responsibility.

## Top 5 Business Doubts and Objections When Implementing Neural Networks

We conducted research on the experience of various companies from different fields and real-life cases, and a conclusion suggests itself: most problems with implementing neural networks are not technical. They are related to expectations, management, and process maturity – everything we've talked about in the book.

- **We'll implement AI because it's trendy.**

The first and most common problem is the lack of a clear goal. Companies often don't understand why they need AI or what exactly it should do. We see this again and again: a business starts implementing neural networks not because there's a specific task, but because "everyone is doing it." As a result, AI appears, but business effect does not.

*Reports from EY and IBM<sup>14</sup> directly state that a significant portion of companies face financial losses precisely due to the lack of a clear implementation strategy, understandable metrics, and clear*

---

<sup>14</sup> <https://www.netshell.co.uk/resources/news-and-insights/blog/top-5-business-challenges-with-ai-adoption/>

*responsibilities. AI is launched as an experiment, and it remains just that – an expensive one.*

In large companies AI initiatives are often shut down or frozen after just a few months: they expected "automation of everything" but got a set of disparate pilots without a clear outcome. For example, the Manipal Hospital network of private hospitals, located in many Indian states, implemented the Watson for Oncology system (incidentally, developed by IBM) in 2016, which was supposed to help doctors analyze medical data and make prescriptions. In the end the project was terminated in 2018. The system did not meet expectations – its conclusions were not sufficiently useful for clinical practice, which, as is known, carries huge legal risks.

- **We have little data.**

And here the business is absolutely right. We've also seen this in projects and research: AI doesn't fix data problems; it amplifies them. If data is dirty or incomplete, the neural network confidently produces incorrect conclusions.

*IBM<sup>15</sup> and Stack AI in their reviews directly point out that data quality and the risk of bias are one of the main fears of companies when implementing neural networks. Energy and infrastructure companies that tried to use AI for forecasts and equipment maintenance encountered the fact that data from old systems simply didn't match. Until they brought order to the sources, AI was useless, sometimes even harmful. Returning again to IBM's developments, they long promoted their Watson AI project as a revolutionary product for medicine. The project cost about \$4 billion, but it turned out the model didn't work with real, "dirty" medical data, didn't fit well into workflows, and produced unclear results. In the end, IBM had to sell part of the Watson Health assets for just \$1 billion and significantly downsize the division.*

- **Our team lacks people who know how to work with this.**

We hear this third objection especially often from managers: Who will manage all this? You need to know how to set tasks for AI, check results, and integrate it

---

<sup>15</sup> <https://www.ibm.com/think/insights/ai-adoption-challenges>  
<https://www.stack-ai.com/blog/the-biggest-ai-adoption-challenges>

into processes. Such skills are often simply absent in teams, and without training and a clear understanding of the implementation process, they won't appear out of nowhere.

Most companies acknowledge a skills gap. And not only technical but also managerial ones. People either fear AI or use it too superficially. It turns out the tool exists, but the effect doesn't.

Only one third of companies systematically train employees to work with AI, be it internal training or outsourcing, even if they are already implementing it into processes.

*In 2023, a loud and telling case happened at Deloitte Australia – one of the world's largest consulting brands, not some "garage company." Deloitte's Australian division was preparing a large analytical report for a government client, several hundred pages long, with a contract budget of about AUD 440,000 (approximately USD 290,000). During preparation, the team officially used generative AI based on Azure OpenAI.*

*Problems began when an independent researcher started reading the document carefully. The report contained errors all too familiar to anyone who has worked with neural networks: non-existent sources, incorrect references, fabricated or distorted quotes.*

*These were precisely those hallucinations usually warned about in theory but often ignored in practice.*

*What became the main stumbling block? These errors went through the entire internal control process! Besides the quality of the government contract report, trust in the overall process was damaged: if such a document can contain fabricated sources, what else wasn't checked?*

*The company acknowledged the problem, publicly explained what went wrong, and agreed to return part of the funds to the client. This case well illustrates reality: the error occurred at the level of people and processes. AI was used as a substitute for attentiveness and professional responsibility. Where "Human in the Loop" was needed, the human turned out to be "out of the loop."*

- **We can't calculate real benefits and ROI.**

Even where AI seems to work, businesses often don't understand: has it actually gotten better? Many companies, in fact, don't know how to calculate ROI from AI projects, especially when it comes to quality and risk reduction.

More than 70% of organizations in the US market lack clear AI effectiveness metrics, even when projects are already scaled. They think AI is implemented, reports

are done faster, texts are more literate, but no one can really say how this affected profit or decisions. As a result, initiatives "hang in the air" and are the first to face budget cuts.

- **Employee Resistance**

And finally the most underestimated problem – the human factor. According to BCG, up to 70% of difficulties in implementing AI are related not, as one might expect, to technology, but to changes in people's work<sup>16</sup>.

We all know this perfectly well: employees don't trust the system, ignore its prompts, or fear that AI is "monitoring" or "will replace" them. Companies implemented chatbots in support, but call center employees continued to work the old way. Clients received strange answers, trust fell, and the project had to be redone almost from scratch.

---

<sup>16</sup> [https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value?utm\\_source=chatgpt.com](https://www.bcg.com/press/24october2024-ai-adoption-in-2024-74-of-companies-struggle-to-achieve-and-scale-value?utm_source=chatgpt.com)

## Top 5 Opportunities and Trends (Despite the Fears and Failed Cases Above)

- **Clean Data First, Then AI**

The first and perhaps main idea for implementing neural networks in your work is a sustained trend which can also be viewed as an opportunity. Businesses that succeeded in implementation started not with models, but with data. First, you need to bring order to sources, processes, and access, and only then connect neural networks. The case below is a living example that implementation works only where order already exists. If you want to use neural networks for forecasts and optimization, you'll first have to invest in data and processes, and only then in models.

*For example, the British oil and gas company Shell implemented neural networks for predicting equipment failures and optimizing production. But the key point is that Shell invested hundreds of millions of dollars over the years in digital infrastructure, sensors, data analysis platforms, and digital twins. In this case, AI is seen as a superstructure over a mature and established system, which was thoroughly prepared before implementation.*

- **AI – Your Assistant, Not an Autonomous Employee**

Another important opportunity in implementing neural networks in your work is abandoning the template of “let AI do everything itself.” We see that successful companies use neural networks as assistants or a kind of interns. AI can help you, for instance, search for information or analyze data, but the final verification of the reliability of this information and the correctness of the analysis remains with the human.

*The American financial conglomerate Morgan Stanley in 2023 implemented generative AI as an internal assistant for its consultants. The model helps quickly find information in thousands of internal documents and recommendations but does not communicate directly with clients.*

*Result: increased consultant productivity and reduced errors because AI is embedded in the workflow, not living its own life. This is exactly the conclusion we've drawn more than once from the cases considered in the book: AI does not replace human thinking; it enhances it.*

- **Focus on Specific Tasks**

The third idea is abandoning large, vague initiatives in favor of narrow and familiar scenarios. Companies have

learned to stop trying to implement AI everywhere possible and started automating individual processes.

UPS uses AI to optimize delivery routes. The ORION system (On-Road Integrated Optimization and Navigation) is a clearly limited tool that saves fuel and time. The company officially announced savings of hundreds of millions of dollars from this approach:

- Savings of 100+ million miles of travel per year.
- Reduction in fuel consumption by approximately 10 million gallons per year.
- Savings of \$300–400 million annually.<sup>17</sup>

- **Local Models and Closed Data Loops**

It's worth mentioning data confidentiality separately, as it is a direct and genuine business risk. Any company has information that should not leave its boundaries. When such data enters public services, the company essentially loses control over it. Local models and closed loops are needed precisely to preserve this control and not turn convenience into a problem. This is

---

<sup>17</sup>

<https://about.ups.com/us/en/our-stories/innovation-driven/orion.html>

a separate reason why corporations and the public sector actively choose this approach.

Local or private LLMs allow you to strictly set rules: what data the AI works with, who has access, what can be uploaded, and what is prohibited. This reduces the risk of intellectual property leakage, simplifies compliance with security service requirements, and removes employee fear that they are being monitored or might accidentally do something critical. Then AI implementation becomes not a source of anxiety but a normal working tool with clear boundaries. And this is precisely the basis for the sustainable and conscious implementation of any technology.

- **Change Management and People Training**

And here we come to the most important point, often ignored. The least but not last. This point correlates very clearly with the last point in the previous block about doubts and objections. Top management is surprised: "We bought the tool, but employees aren't using it." But honestly, there's nothing surprising about it.

*PwC understood quite early a simple thing: if you just give people an AI tool, it will either be used superficially*

*or ignored altogether, or become a source of errors and contradictions. Therefore, they focused not only on systematic training of their employees and conducted a series of trainings.*

*And for the thousandth time, we return to the same idea – AI was not positioned as a replacement for employees; it was presented as a tool to help them, which in turn saves time by assisting with routine tasks and providing hints, but does not remove responsibility. As a result, resistance decreases not because people were forced, but because it became clear to them why it's needed at all.*

Any objections from business are a perfectly obvious reaction to the unknown. But you must understand that these are not the only doubts and illusions about neural networks and, naturally, not the only steps that will make your business soar to immeasurable heights. Each case is individual, each company follows its own path, and all this, once again, is the result of cooperation, effective communication, experience, the ability to build strategy, and make the right decisions by real, living human employees!

## Conscious Use

Therefore, the key question is not "to use or not to use AI," but **how** to do it. Consciously – means in a way that you control the neural network, not it your business. Not trying to patch a leaky bucket with neuro-avatars and agents, but automating what already works clearly and stably without them.

In all these processes, the principle of Human in the Loop remains critically important: the human must be in the loop – checking, correcting, making final decisions. Neural networks are excellent at helping to accelerate, but they are poor at handling responsibility.

If you look more broadly, nothing fundamentally new has happened here. Even before, 95% of websites, applications, and marketing materials were mediocre. Now they've simply added plastic, smooth text generated by neural networks or lazy marketers. And success is still achieved by those who think, combine data, experience, and intuition, manage teams, and do not shift responsibility to tools or "magical platforms." Those who stop thinking for themselves will remain stagnant – regardless of how many AI agents and automations they have.

I want to say something important to leaders separately. Don't worry: no "youth with AI" will push you aside anywhere. You have what cannot be "vibe-coded" – experience, connections, the ability to work with people, to feel context and the real world. A student loaded with agents can quickly generate texts and prototypes, but they often don't know how to contact people, don't sense risks, and don't understand consequences. And it's precisely now that for adults, flexible, non-fossilized professionals, a second youth is beginning – that very startup period when hypotheses can be tested quickly and cheaply, and technology finally starts playing to your advantage.

You can acquire the technological skill – learn to work with neural networks as confidently as you work with people. But your brain, your intuition, and accumulated experience – this is the main competitive advantage that won't go anywhere. Develop it, maintain mental plasticity, and don't forget a simple thing: we live in a physical world. We work with people, for people, and because of people. And here, they still buy from people.

# What's Next?

In summary, putting it all together we see a clear picture:

<b>Common Doubts and Implementation Errors</b>	<b>Opportunities and Solutions</b>
No clear goal and specific tasks	One useful scenario is better than ten pilots for show
Dirty and unprepared data	Prioritize data and processes
Untrained team	Position AI as an assistant
Effect and result not measured	Establish comprehensive and clear metrics
Resistance to change from employees	Training and answering the "why" question for employees
Either use AI or not, no third option	The key question is not "to use or not," but how exactly to use
AI will do everything itself	The business controls the neural network, not vice versa
Implementation will solve bad process problems	AI implementation reveals bad processes
The main thing is to choose the right model	The main thing is to choose the right context of application
The more agents, the better	One stable scenario is better than many non-working solutions
Implementing AI is an IT task	Implementing AI is a management task
Neural networks will soon replace all of us	They will replace everyone who can think + use neural networks

We have done, in fact, a very mature thing: looked at AI without rose-coloured glasses. And if after reading this book you have more clarity and a little less belief in magic in your head – then we didn't undertake all this in vain.

We were happy to share with you what we've learned through practice. Now you understand where neural networks help, where they hinder, where they require

control, and where it's better not to touch them at all until there is order in the foundation. I truly hope that now your perspective has become a little wider and your thinking more flexible and prepared for what we might call the new technological paradigm. And you approach decisions more calmly – with the question "why" and with the responsibility and understanding that acceleration without a steering wheel is a direct path to the cliff.

The goal of this book is by no means a call to urgently implement AI everywhere possible; it's more food for thought so that you can decide for yourself whether you're ready to use neural networks now. Or maybe, on the contrary, you want to steer clear of this "devil in the computer" and continue working with traditional methods. Both options are normal. The decision is always yours, and that, by the way, is the healthiest outcome.

If we were to add one final thought from us, it would probably be this: neural networks in business are used not to become faster and better than everyone else, but still to become more honest with yourself. They very quickly show where your company has a clear process and where there's fog; where there is meaning and responsibility, and where there is hope for a miracle.

And if you take away from this part at least one thing – the habit of thinking, checking, and keeping the human in the loop – then you are already a step ahead of the majority. And we will be sincerely glad if this step proves useful for you – in business, in analytics, in management, even in correspondence with your own team.

If you want to continue this journey further – to understand more deeply, implement approaches in your company, discuss cases, or work together – we have neural network training, individual consultations, and a channel where we regularly show live projects from our team: vibe-code solutions, experiments, and just good news from the world of artificial intelligence.

We also hold conferences, corporate training, and help businesses build work with AI without empty promises and unnecessary noise.

Contact us – and my team and I will select the best solution for you individually.

