

Thoughtware Designing in the Intelligence Age

A philosophy of product planning
and design for an age in which
software begins to think.

Achin Kumar

Founder, Session Hero

Co-founder, Roro.



thoughtware.

CONTENTS

Preface · Why I Wrote This Book 1

A personal reflection on why the Intelligence Age demands a new philosophy of product design.

Introduction · We Are Entering the Intelligence Age 5

How humanity transitions from one age to another, why intelligence is the defining capability of our time, and why a new philosophy of design is needed.

PART I · THE SHIFT

Chapter 1 · The Intelligence Age 14

Why history is shaped by new capabilities, how language enabled general intelligence, and why Large Language Models mark the beginning of the Intelligence Age.

Chapter 2 · The Rise of Thoughtware 22

The evolution of software from deterministic programs to machine learning, language models, and autonomous agents, and what fundamentally changed.

Chapter 3 · Consumers No Longer Operate Software 38

How consumer psychology is evolving from operating software to directing intelligence, and why expectations are permanently shifting.

Chapter 4 · The Intelligence Product Leader 46

Why product managers, designers, researchers, and engineers must rethink their craft, and the new disciplines they need to master, from LLMs and memory to retrieval, agents, evaluation, and intelligence architecture.

PART II · THE PHILOSOPHY

THEME I · REDUCING COGNITIVE LOAD**Chapter 5 · Low-Cognitive First Draft 62**

Why consumers increasingly expect software to eliminate the blank page and begin with thoughtful work instead of empty states.

Chapter 6 · Point & Fix 69

Why consumers increasingly expect surgical editing instead of repetitive workflows, and how pointing has become the new command language.

Chapter 7 · Intent over Instructions 79

Why consumers now express outcomes instead of procedures, and why products must infer the “how” from the “what.”

Chapter 8 · Conversation over Configuration 89

Why natural language is replacing menus, settings, and forms as the primary interface for expressing intent.

THEME II · SOFTWARE THAT KNOWS YOU**Chapter 9 · Context is a Feature 101**

Why context is no longer supporting information but a core product capability.

Chapter 10 · Memory is Part of UX 111

Why every interaction should improve future interactions, making software increasingly personal rather than stateless.

THEME III · INTELLIGENCE AS THE INTERFACE**Chapter 11 · Recommend, Don’t Respond 123**

Why intelligent products should guide decisions rather than merely answer questions.

Chapter 12 · The Devil is in the TL;DR 132

Why intelligence is measured not by how much it generates, but by how effectively it reduces cognitive effort.

Chapter 13 · Visible Thinking 141

Why consumers increasingly trust software that reveals its reasoning instead of hiding it.

Chapter 14 · Intelligence over Aesthetics 150

Why intelligence is becoming the strongest driver of product delight, with visual polish becoming expected rather than differentiating.

THEME IV · A NEW INTERACTION MODEL**Chapter 15 · Language over Layout** 160

Why language is becoming the universal interaction layer across text, voice, images, documents, interfaces, and physical environments.

Chapter 16 · Co-pilot the Decision 169

Why the future of software lies in collaborative decision-making, balancing human judgment with machine intelligence.

PART III · APPLYING THE PHILOSOPHY**Chapter 17 · Designing Products for the Intelligence Age** 182

A practical framework for applying the principles to product strategy, system architecture, user experience, workflows, and organizational design.

Chapter 18 · Beyond the Interface 189

Where intelligent products are heading, how consumer expectations will continue to evolve, and why the Intelligence Age is only just beginning.

Glossary 196

A quick reference to the ecosystem of thoughtware: every concept and technique in the book, with chapter and page.

Preface

Why I Wrote This Book

Every few decades, a technology comes along that does more than improve the way we work. It quietly changes the way we think. I believe we are living through one of those moments. Over the past few years, I have had the opportunity to design and build products powered by Large Language Models and agentic AI. Like many others, I initially viewed AI as another technological advancement to be integrated into products, a remarkably capable one, but still another capability in the long evolution of software.

The more products I built, however, the more I realized that I was asking the wrong questions. I was trying to understand what AI could do. The more important question was what AI was doing to the people using it. I began noticing subtle but consistent changes in behaviour. Users no longer wanted to master complex interfaces before becoming productive. They expected software to understand their intent. They were increasingly comfortable reviewing an AI-generated first draft instead of starting from a blank page. They expected software to remember previous conversa-

tions, preserve context, make recommendations, and remove routine cognitive effort from their work.

The products that delighted people the most were rarely the ones with the most polished interfaces. They were the ones that behaved intelligently. That observation stayed with me. For decades, product design has focused on helping people operate software more effectively. We invested enormous effort into making interfaces clearer, workflows simpler, navigation easier, and interactions more intuitive. Those efforts transformed the software industry, and they remain just as important today.

But they were developed for a world in which software was fundamentally deterministic. Software executed instructions, while people carried the responsibility of understanding problems, making decisions, remembering context, and navigating complexity. That relationship is beginning to change. Software is no longer limited to executing instructions. Increasingly, it can understand language, reason through ambiguity, retrieve knowledge, generate alternatives, make recommendations, and perform meaningful cognitive work alongside its users.

The interface is no longer the only thing being designed. Intelligence has become a design material in its own right. As I spoke with designers, founders, product managers, and engineers, I noticed something else. While everyone recognized that AI was important, conversations often revolved around features rather than philosophy. Teams discussed prompts, models, and copilots, but there was far less discussion about how thoughtware fundamentally changes consumer expectations and, consequently, the principles by which products should be designed.

This book grew out of that gap. It is not intended to be a book about artificial intelligence. Nor is it a book about prompt engineering, model architectures, or the latest technological breakthroughs. Those topics evolve far too quickly to anchor a lasting design philosophy. Instead, this book is about people. It is about how technology changes behaviour, how changing behaviour reshapes expectations, and how those expectations demand an entirely new way of thinking about product design.

The twelve principles presented in this book are not intended to replace the foundations of user experience design. Good visual design, accessibility, usability, performance, information architecture, and human-centred design remain indispensable. Every new age is built upon the one before it. The Intelligence Age simply asks us to build another floor. That new floor is built from memory instead of state, intent instead of workflows, reasoning instead of automation, collaboration instead of execution, and intelligence instead of interface alone.

I do not claim that these principles are complete, nor that they represent the only way to design intelligent products. We are still in the earliest days of this transition. Many ideas in this book will undoubtedly evolve as the technology matures and as consumers develop new expectations of their own. If this book succeeds, it will not be because every prediction proves correct.

It will be because it gives product builders a common language for discussing the challenges ahead. Every age eventually develops its own philosophy of design. The Industrial Age shaped industrial design. The Information Age shaped interaction design and modern user experience. The Intelligence Age deserves a philo-

sophy of its own. This book is my attempt to begin that conversation. Part of that language is a word I have found myself needing: thoughtware, software that thinks rather than software that executes.

Introduction

We Are Entering the Intelligence Age

Every generation believes it is living through extraordinary technological change. Most are right. But only a handful of generations experience a breakthrough so fundamental that it changes not only the tools people use, but the way they think, work, and behave. Those moments don't simply create better technology. They redefine what people expect from the world around them.

I believe we are living through one of those moments. As we move from the Information Age to the Intelligence Age, the source of consumer delight is rapidly shifting, from beautiful interfaces and abundant information to software that is intelligent, contextual, personalized, and capable of meaningful reasoning. Today's consumers don't simply expect software to execute commands. Increasingly, they expect it to understand them, remember them, guide them, and remove routine cognitive effort from their lives.

This book is about designing products for those expectations.

Every New Age Is Built on the Previous One

A new age never replaces the one before it. It builds upon it. History is less like a sequence of replacements and more like the construction of a tower. Every generation adds another floor. The previous floors do not disappear. They become the foundation. Electricity did not replace agriculture. The internet did not replace electricity. Smartphones did not replace the internet. Each breakthrough became infrastructure upon which the next breakthrough was built.

The same applies to product design. The Information Age gave us remarkable principles for building digital products. Visual hierarchy, usability, accessibility, information architecture, interaction design, responsive layouts, and human-centered design transformed software from something only specialists could use into something billions of people could use effortlessly. Those principles remain indispensable. But they are no longer sufficient. Not because they are wrong, but because the world they were created for has changed.

For decades, software behaved exactly as it was programmed to behave. Today, software increasingly participates in solving problems alongside its users. The next floor of the tower is intelligence. Building products for that floor requires another layer of design thinking.

Consumer Psychology Is Already Changing

Every significant technological breakthrough changes consumer behaviour long before it changes industries. The first people to use automobiles expected to travel farther. The first people to use smartphones expected information to be available everywhere. Those expectations quickly became normal. The same process is happening with thoughtware. Consumers are increasingly expecting software to draft before they write, summarize before they read, recommend before they search, organize before they plan, and answer questions they haven't fully articulated yet.

They are spending less time creating from scratch and more time reviewing, refining, and directing. The blank page is quietly disappearing. This is not merely a productivity improvement. It represents a deeper change in the relationship between humans and software. For decades, users accepted that software required them to think first and operate second. They learned workflows, navigated interfaces, remembered where features lived, and translated their intentions into commands software could understand.

Increasingly, that burden is shifting. Consumers are beginning to expect software to understand their intent instead.

From Operating Software to Directing Intelligence

One observation kept appearing as I worked on intelligent products. Consumers no longer wanted to operate software. They wanted to direct intelligence. That difference may appear subtle, but it fundamentally changes what great product design looks like. Operating software requires interfaces. Directing intelligence requires collaboration. Traditional software waited for in-

structions. Thoughtware increasingly participates in solving problems. It asks clarifying questions.

It proposes alternatives. It retrieves relevant knowledge. It drafts. It critiques. It remembers. It recommends. Instead of acting only after every explicit command, it increasingly shares responsibility for the work itself. The relationship between humans and software is beginning to resemble collaboration rather than operation.

More Than Another Technology Shift

Whenever a major technological breakthrough arrives, our instinct is to focus on the technology itself. History suggests we should instead focus on the people. Steam engines changed transportation, but more importantly they changed where people lived, how factories operated, and how economies grew. The internet changed communication, but more importantly it changed how people discovered information, built relationships, learned new skills, and made decisions.

Large Language Models will undoubtedly transform software. But the larger transformation will occur in consumer psychology. The products that define this era will not simply contain artificial intelligence. They will be built around entirely different expectations of how software should behave. Software that remembers instead of forgetting. Software that understands instead of asking. Software that recommends instead of merely responding.

Software that collaborates instead of waiting. Software that increasingly adapts itself to people instead of asking people to adapt to it.

A New Philosophy of Design

Every age eventually develops its own philosophy of design. The Industrial Age shaped industrial design. The Information Age shaped interaction design and modern user experience. The Intelligence Age demands another evolution. Not because aesthetics have become unimportant. Not because usability no longer matters. But because intelligence has become a first-class design material. Just as designers once learned to shape pixels, screens, interactions, and workflows, we must now learn to shape context, memory, reasoning, recommendations, autonomy, and collaboration.

The questions product teams ask are already changing. How much should the system remember? When should it recommend instead of asking? When should it act autonomously? How should it reveal its reasoning? How should humans and intelligence make decisions together? These are no longer engineering questions alone. They are product questions. They are design questions. They are experience questions.


This Book

This book is an attempt to answer those questions. Not by cataloging AI features or interface patterns, but by proposing a philosophy for designing products in the Intelligence Age. The principles presented in the chapters ahead are rooted in one simple observation:



Consumers are transitioning from operating software to directing intelligence.

In the Information Age, users adapted themselves to software. In the Intelligence Age, software increasingly adapts itself to users. That behavioural shift changes almost every assumption we have held about great product design. The chapters that follow explore why this shift is happening, how it changes software and consumer behaviour, how it reshapes the role of product builders, and finally, the principles that can guide us in designing products worthy of the Intelligence Age.



PART I

The Shift

Every age of humanity has been defined by a capability that fundamentally expanded what people could do. When that capability became abundant, it didn't simply improve existing ways of living. It reshaped civilization itself. New industries emerged, old ones disappeared, jobs evolved, economies reorganized, and entirely new expectations of the world were formed. Technology rarely transforms society directly. It transforms people first. People adopt new behaviours, those behaviours become habits, habits become expectations, and those expectations eventually reshape every product, service, institution, and business built around them.

This pattern repeats throughout history. The automobile changed far more than transportation. It changed the design of cities, commerce, tourism, and even the way people planned their lives. The internet did far more than connect computers. It changed how knowledge is discovered, relationships are formed, businesses are created, and trust is established. Smartphones did more than place computers in our pockets. They fundamentally changed our expectations around convenience, responsiveness, and immediacy.

Every technological revolution quietly rewrites what consumers consider normal. The emergence of Large Language Models and intelligent agents represents another such shift. It is more than a software breakthrough. It changes how humans perform cognitive work itself. For the first time in history, intelligence is becoming abundant, accessible, and programmable. That changes the role of software. It changes the role of people. And ultimately, it changes the principles by which products should be designed.

Before exploring those principles, we must first understand the shift itself. What defines an age? Why does intelligence represent a new capability rather than another incremental improvement? How does that capability reshape software? How does software reshape consumer behaviour? And why must product builders rethink their craft? The chapters in this section answer those questions.

The Intelligence Age

Ages Are Defined by Capabilities

History is often taught as a sequence of civilizations, empires, and dates. A more useful way to understand it is as a sequence of capabilities. Humanity does not enter a new age because a calendar changes. It enters a new age because it acquires a capability that fundamentally changes what ordinary people can do. Once that capability becomes widely accessible, it alters behaviour. New behaviours become expectations, and those expectations eventually reshape civilization. (See Fig 1.)

The names we give to historical ages reflect this pattern. The Stone Age was not defined by the existence of stones. Stones had always existed. What changed was humanity's ability to shape them into tools. Those tools transformed hunting, construction, protection, and survival. The Bronze Age was not simply about discovering bronze. It was about learning to manipulate metals into stronger tools, weapons, and infrastructure. This capability enabled larger settlements, more sophisticated trade, and increasingly organized societies.


The Iron Age expanded those possibilities even further. Agriculture improved, military power shifted, and civilizations grew larger and more complex. Thousands of years later, the Industrial Revolution multiplied human physical capability. For the first time, machines could perform physical work at a scale impossible for individuals. Factories emerged, cities expanded, and entire economies reorganized themselves around industrial production.

The twentieth century introduced another transformation. Computers converted physical information into digital information. Documents became files, photographs became pixels, music became bits, and communication became electronic. The internet built upon that foundation by making information globally accessible. Knowledge became searchable, communication became instantaneous, and geography became dramatically less important than it had been for centuries. Each age introduced a new capability. None of these technologies simply improved existing systems. They fundamentally changed how people lived and worked.

The question then becomes: what capability defines our age?

The four ages and what each made abundant

FIG 1



From Information to Intelligence

The defining capability of the last few decades was information. Computers allowed us to store information. The internet allowed us to distribute it. Search engines allowed us to retrieve it almost instantly. Smartphones made it permanently accessible. The defining challenge of the Information Age was no longer obtaining knowledge but navigating an overwhelming abundance of it. Today, another capability is emerging. (See Fig 2.)


Intelligence. Not human intelligence. Not consciousness. Not artificial general intelligence. Rather, the ability for software to participate meaningfully in cognitive work. For decades, software excelled at executing explicit instructions. If a developer anticipated every possible situation, software behaved beautifully. If reality fell outside those assumptions, software failed. Traditional soft-

ware was deterministic. Machine learning represented a significant step forward. Rather than explicitly programming every rule, we trained systems to recognize patterns. Spam filters improved. Recommendation engines became practical. Fraud detection advanced. Medical diagnosis became more accurate.

These systems were remarkable, but they remained specialists. Each model solved one category of problems. Each required carefully prepared data. Each operated within relatively narrow boundaries. General reasoning remained uniquely human. Large Language Models fundamentally changed that trajectory.

When a capability becomes abundant,
something else becomes scarce.

FIG 2



Why Language Matters

What makes Large Language Models different is not simply that they generate text. It is that they operate on language itself. Language is often thought of as a communication system. In reality, it is humanity's greatest abstraction. It allows us to compress complex ideas into symbols and manipulate those symbols without directly experiencing reality. Consider planning a business, designing a building, negotiating a contract, writing a novel, or imagining a future that has never existed. None of these activities require the physical object to exist. They happen almost entirely inside symbolic representations created through language.

Language allows humans to think beyond immediate experience. It enables imagination, planning, reasoning, teaching, and collaboration. Civilization itself depends upon this abstraction. There is a simple thought experiment that illustrates this surprisingly well. Try explaining that there is a cup of coffee sitting on a wooden table without using the words cup , coffee , table , wood , or legs .

Most people immediately begin to struggle. The difficulty isn't merely finding alternative words. The difficulty is organizing the thought itself. Your understanding of the world has been encoded through language. Remove enough of that abstraction and thinking itself becomes difficult. Language is not simply how humans communicate. It is one of the primary ways humans think.


Evolution Took the Same Path

Nature followed a remarkably similar journey. The earliest forms of intelligence consisted of tiny collections of neurons producing equally simple behaviours: move toward food, move away from danger, avoid obstacles. Simple inputs produced simple outputs. Over millions of years, intelligence became increasingly specialized. The Venus flytrap is a fascinating example. Tiny sensory hairs detect pressure, and when enough stimulation occurs within a short period, the trap closes almost instantly before beginning digestion. It is an extraordinary behaviour, but one designed to solve exactly one problem. (See Fig 3.)

Classical software resembles the earliest stage of intelligence. Traditional machine learning resembles the second. Both are highly capable within carefully defined boundaries. General intelligence required another leap. It required abstraction. Language became that abstraction.

FIG 3

Language as the abstraction layer



The Emergence of Programmable Intelligence

This is what makes Large Language Models fundamentally different from previous generations of software. Rather than manipulating numbers alone, they manipulate the same symbolic system humans use for reasoning. That does not mean today's models think exactly like humans, nor does it imply consciousness. Those are different questions altogether. What matters is that software can now participate in kinds of cognitive work that previously required human minds. It can draft documents, summarize information, brainstorm ideas, compare alternatives, critique proposals, generate plans, explain unfamiliar concepts, and reason across large bodies of knowledge.

For thousands of years, intelligence was constrained by biology. If you wanted another intelligent perspective, you needed another human being. Today, intelligence has become available through software. Not perfect intelligence. Not human intelligence. But programmable intelligence. Throughout history, every age has democratized a capability that was once scarce. The Industrial Age democratized physical power. The Information Age democratized access to knowledge. The Intelligence Age is beginning to democratize cognitive capability itself.

When physical labour became abundant, work changed. When information became abundant, knowledge work changed. As intelligence becomes abundant, cognitive work will inevitably change as well. This is why Large Language Models should not be viewed as merely another software technology. They represent the emergence of a new capability upon which entirely new products, professions, business models, and consumer expectations will be built.

Every age eventually develops its own architecture, its own institutions, and its own philosophy of design. The Intelligence Age will be no different. The question is no longer whether thoughtware will reshape the products we build. It already is. The real question is whether we are prepared to design for a world where intelligence is no longer scarce.

The Rise of Thoughtware

Software has always existed to amplify human capability. What has changed over the past seventy years is not merely the amount of computation available to us, but what computation is expected to do. Every major breakthrough in software has shifted another layer of cognitive effort away from humans and into machines. Large Language Models represent the largest shift in that progression so far.

To understand why, we need to look at how software itself has evolved.

The Era of Deterministic Software

For most of computing history, software behaved exactly as instructed. Every possible action had to be anticipated by its creators. Every workflow was explicitly designed. Every menu, button, dialog, and error message represented a decision someone had already made on behalf of the user. Software never improvised. It never inferred. It never understood intent. It simply executed instructions.

This fundamental limitation shaped the entire discipline of product design. Since software could not understand users, users had to understand software. They learned menus. They memorized shortcuts. They navigated workflows. They translated their intentions into the exact sequence of actions the application expected. The burden of understanding rested almost entirely on the human. Many of the UX principles we rely upon today emerged in response to this reality. Discoverability, consistency, clear navigation, progressive disclosure, visual hierarchy, and affordances all exist because users needed help operating increasingly complex software.

These principles remain as valuable today as they were then. But they were designed for a world in which software executed rather than reasoned.

The Era of Predictive Software

Machine learning changed the relationship between humans and software, but only slightly. Instead of explicitly programming every rule, developers could train models to recognize patterns in data. Spam filters became remarkably accurate. Recommendation systems suggested products and movies. Maps predicted traffic. Keyboards predicted the next word. Search engines became dramatically more relevant. Software became better at making predictions.

Yet the overall relationship between humans and computers remained largely unchanged. Each model solved one carefully defined problem. A recommendation engine recommended. A fraud detector detected fraud. An image classifier classified images. The intelligence was narrow, specialized, and embedded in-

side otherwise deterministic applications. The interface remained the product. Machine learning simply made that interface smarter.

The Emergence of Thoughtware

Large Language Models introduced something fundamentally different. Instead of building separate models for writing, summarization, coding, translation, brainstorming, planning, reasoning, research, and conversation, a single model could perform all of these tasks through the same interface. Language. That is a profound shift. For the first time, software no longer required every interaction to be explicitly designed in advance. Users could simply describe what they wanted, and software could determine how to achieve it. (See Figs 4 and 4.1)

The interface began to move away from controls and toward intent. Instead of asking users to learn software, software increasingly began learning users. That reversal may become one of the defining characteristics of the Intelligence Age.

This is a different category of software, and it deserves a different name. Software executes instructions. Thoughtware reasons toward outcomes. It takes intent rather than commands, holds context rather than state, produces judgment rather than output, and improves with use rather than with releases. The distinction matters because thoughtware cannot be specified the way software is specified, and it cannot be verified the way software is verified. It has to be described, evaluated, and improved. Everything that follows in this book is an attempt to work out what that means for the people who design it.

Three eras of software

FIG 4

A taxonomy for how software thinks and behaves.

As intelligence becomes table stakes, delight shifts from how things look to how intelligently they work.

	 DETERMINISTIC The rules era	 PREDICTIVE The learning era	 THOUGHTWARE The reasoning era
 HOW IT DECIDES	Follows explicit rules. Given the same input, it always produces the same output.	Finds patterns in data. Uses statistical relationships to predict the most likely output.	Reasons over goals and context. Plans, decomposes, and adapts to achieve the best outcome.
 WHAT IT NEEDS TO WORK	Rules and logic. Deterministic code, fixed logic, and structured inputs.	Large amounts of data. Quality data, features, and representations to learn from.	Context, knowledge, and tools. Rich context, memory, feedback, and access to tools and systems.
 HOW IT FAILS	Breaks when rules are wrong or incomplete. Can't handle exceptions outside the rules.	Breaks when data is missing or out of distribution. Struggles with novel situations or noisy, biased data.	Breaks when reasoning goes wrong or context is poor. May hallucinate, take wrong actions, or pursue the wrong goal.
 HOW YOU TEST IT	By checking correctness. Unit tests, assertions, and exact input-output verification.	By measuring accuracy. Train/validation/test sets, metrics (precision, recall, RMSE, etc.).	By evaluating outcomes and behavior. Task success, human judgment, A/B tests, safety, robustness, and continuous evals.

→ **THE SHIFT:** From executing instructions → to predicting outcomes → to reasoning toward goals.

From Tools to Collaborators

Traditional software helped people perform work. Thoughtware increasingly performs work alongside people. The distinction is subtle, but significant. A spreadsheet helps you calculate. A language model helps you think through the calculation. A presentation tool helps you create slides. A language model proposes the narrative, generates the structure, critiques the flow, and suggests improvements. A code editor allows developers to write code.

An intelligent development environment can propose architectures, implement features, explain unfamiliar systems, identify bugs, and recommend better approaches. Software is no longer participating only in execution. It is increasingly participating in

cognition. This changes the role of the human. Instead of performing every cognitive step, people increasingly review, guide, approve, and refine the work performed by intelligent systems.

The machine performs more of the intermediate thinking. The human focuses more on judgment.

From Features to Capabilities

Historically, software products competed by adding features. Every release introduced another screen, another button, another workflow, another configuration panel, or another menu. Products became increasingly capable. They also became increasingly complicated. Users paid for capability with complexity. Thoughtware changes this relationship. A sufficiently capable model can perform hundreds of different tasks through the same conversational interface. Rather than exposing every capability through dedicated workflows, products increasingly expose capability through understanding.

Roadmaps begin shifting away from adding interfaces and toward expanding intelligence. The product grows in capability rather than complexity.

The New Architecture of Software

This transformation is much deeper than the interface. It changes the architecture of products themselves. Traditional software consisted of familiar layers: user interfaces, business logic, databases, and APIs. Every request flowed through deterministic code that produced deterministic outcomes. Intelligent products introduce entirely new architectural layers. Reasoning models determine how problems should be approached. Memory allows systems

to accumulate knowledge about users over time. (See Figs 5 and 5.1)

What thoughtware is

FIG 4.1

Thoughtware is software that understands, adapts, and reasons toward goals. It's a new kind of software built for a new kind of work.

 Traditional software The past: programs we control	 Thoughtware The future: partners that reason
 Takes commands Waits for explicit instructions and executes them.	 Takes intent Understands goals and figures out how to achieve them.
 Holds state Stores data in tables, records, and variables.	 Holds context Maintains meaning, relationships, history, and nuance.
 Produces output Generates answers, reports, or actions.	 Produces judgment Weighs trade-offs and chooses the best path forward.
 Verified Correct if it matches the rules and the tests.	 Evaluated Good if it achieves goals and earns trust.
 Improves with releases Gets better when we ship new versions.	 Improves with use Gets better as it learns and adapts over time.
 THE EMERGENCE OF THOUGHTWARE	
As intelligence moves from automation to augmentation, software evolves from tools we control to partners we collaborate with.	

Retrieval systems dynamically gather relevant information instead of relying solely on training data. Planning systems break large objectives into manageable tasks. Tools allow models to interact with external systems. Skills package procedures the system already knows, so settled work is performed rather than worked out again. Agents coordinate multiple steps toward a goal. Evaluation systems continuously assess confidence, quality, and safety before responses reach the user. These layers are not merely engineering concerns. They directly shape the experience people have with the product.

Whether software remembers too much or too little. Whether it retrieves current information. Whether it asks clarifying questions. Whether it explains its reasoning. Whether it requests approval before acting. These are architectural decisions. They are equally design decisions.

These layers describe what an intelligent product contains. They do not describe what you build. The unit of construction is smaller: a single prompt, with a defined input, a specified model, and an expected output. Call it a cognitive unit. It is the function of the Intelligence Age, and like a function it can be planned, built, tested, evaluated, and improved on its own. Above it sit two further kinds of construct. A skill is a known procedure that performs a capability. An agent owns a goal and works out its own route to it. A cognitive unit decides, a skill performs, an agent pursues. Intelligence is assembled, not summoned.

FIG 5

The component stack of an intelligent product



Skills: The Work Between Decisions and Goals

Between the cognitive unit and the agent there is a gap, and most real systems live inside it. A cognitive unit owns one bounded decision. An agent pursues a goal and decides for itself when it is finished. A great deal of useful work is neither.

Consider an invoice arriving for review. The product extracts the fields, finds the matching order, computes the difference, stops if the difference falls inside the approval threshold, classifies the exception if it does not, assigns a priority, and routes the case. There is real cognition in that sequence. Classifying the exception is a judgement no rule could settle in advance. But there is almost no agency. The product never asks what kind of work it should be doing. It does not reconsider its objective halfway through, or discover a new route because something surprising appeared. Somebody already knew the shape of the work and wrote that shape down.

Calling that an agent confuses intelligence with agency, which are different things. Calling it a pipeline describes how the work happens rather than what work the product knows how to do. The name for the second of those is a skill: a reusable, bounded procedure combining cognitive units, ordinary code and tools to perform a known capability. Its path may branch. It may contain judgement. What distinguishes it is that the shape of the work is largely settled before it runs.

A cognitive unit decides. A skill performs. An agent pursues. These are responsibilities, not sizes.

The temptation is to read these as small, medium and large. That reading is wrong. A skill holding twenty operations may be simpler than an agent holding three. What separates them is what each is responsible for. A decision is answered. A capability is performed. A goal is pursued. The grammar is the tell.

A skill may contain judgement without becoming an agent. A skill that prepares a variance review might ask one cognitive unit which discrepancies are material and another how the policy applies. Both are open decisions. The procedure is nevertheless stable: materiality is always judged after the discrepancy exists, the policy interpretation always receives the applicable policy, and the result always plays the same role in the larger system. The cognition happens inside the procedure. In an agent, cognition helps determine the procedure itself.

This matters for a practical reason. Acting on the world is far too weak a test for agency; ordinary software has sent messages, moved money and opened doors for seventy years. Agency appears when control has been delegated, when the product decides what to attempt next in light of what has already happened, and when finishing is something it has to judge rather than simply arrive at. A skill ends because its procedure ends.

Nor is this a strict hierarchy. It is tempting to draw thoughtware containing agents containing skills containing cognitive units.

That is a useful first sketch and a poor theory. An agent can call a cognitive unit directly. A skill can call another skill. A product can invoke a skill with no agent anywhere in between, and often should. If somebody presses a button marked generate my grocery list and generating that list is a known procedure, waking an agent so that it can decide to perform the procedure adds no intelligence at all. It adds another place for the product to vary.

The same skill may later appear inside something larger, when the request becomes sort out next week's food and use whatever is in the fridge. Now the product owns a goal, and the grocery list is one known step on a path it did not know at the outset. The skill did not change. Its caller did.

A useful skill needs the same things a good product component always needed, named honestly. A name describing a capability rather than a decision. Inputs and outputs, so a caller need not understand the sequence inside. A legible set of things it may do. Visible dependencies. Where it can act, an explicit authority, because a skill that prepares a payment differs from one that submits it in the way that matters most. And failure behaviour worth the name, since a missing record, an unavailable integration, an abstention and an exhausted retry do not mean the same thing and should not collapse into one generic error.

When a Path Becomes a Skill

There is a second reason the middle layer matters, and it reaches further than the first. An agent meeting a new class of problem investigates, tries routes, checks results and recovers from mistakes. After enough cases a pattern emerges. For one kind of exception it almost always retrieves the same three documents,

asks the same judgements in the same order, applies the same check and prepares the same package. What was discovered case by case has become a known way of doing the work.


At that point continuing to deliberate is waste. The path can be named and kept as a skill. The agent still owns the larger goal; it has simply acquired a capability it no longer has to reconstruct. This is the same movement by which a decision that has become regular enough turns into ordinary code, one level up. A decision closes when judgement stops being needed. A path stabilises when orchestration stops being needed. Both take uncertainty out of the moment of use and return it to design.

It follows that a maturing product should not simply acquire better agents. Parts of its agentic behaviour should disappear, and its agents should get smaller, because the product has learned what it is doing. Autonomy is valuable exactly where procedure ends. Using it everywhere does not make a product more intelligent. It makes settled work uncertain again.

Three responsibilities

A cognitive unit decides, a skill performs, an agent pursues

FIG 5.2



The Rise of Agentic Software


Language models represent only the first phase of this transition. Today's intelligent systems primarily respond. Tomorrow's systems will increasingly act. An agent does more than answer a question. It understands an objective, creates a plan, selects the appropriate tools, executes multiple steps, evaluates the outcome, recovers from failures, and continues working until the objective has been achieved, or until human judgment is required. (See Fig 6.)

The interaction changes from requesting outputs to delegating outcomes. This is perhaps the most important transition still ahead. Products will increasingly be judged not by the quality of their responses, but by the quality of the work they complete.

The cognitive unit

FIG 5.1


The cognitive unit (CU) is the new building block of software.
It turns inputs and goals into intelligent action.



Trust, transparency, autonomy, memory, recovery, and collaboration will become first-class design concerns.

One response vs the agentic loop

FIG 6



Designing Thoughtware

Every era of software has required a different philosophy of design. The desktop era emphasized commands. The web emphasized navigation. Mobile emphasized simplicity. The Intelligence Age emphasizes collaboration. Products are no longer collections of screens connected by workflows. They are becoming systems that understand, remember, retrieve knowledge, reason, recommend, and act. As software evolves, so must the people who build it.

Understanding interfaces alone is no longer enough. Product builders must now understand intelligence itself. Because consumers no longer evaluate software solely by what it can do. They increasingly evaluate it by how intelligently it behaves while doing it. That evolution in software is only one side of the story. The

other side is even more important. As software changes, consumers change with it.

The next chapter explores how this new generation of intelligent products is reshaping human expectations, and why consumers are gradually transitioning from operating software to directing intelligence.

Consumers No Longer Operate Software

Technology does not change consumer expectations overnight. It changes behaviour first. The first people who used search engines still browsed directories. The first people who bought smart-phones still carried paper maps. The first people who adopted cloud storage continued saving files locally. Habits rarely disappear the moment a new technology arrives. But once people experience a fundamentally better way of doing something, they rarely go back.

The old behaviour slowly begins to feel unnecessary. Eventually, it feels primitive. That is precisely where we find ourselves today. The emergence of intelligent software is changing not only what people can accomplish, but what they expect software to do on their behalf. This chapter is not about technology. It is about psychology.

Expectations Are Always Relative

Consumers rarely evaluate products in isolation. They compare every experience against the best experience they have had anywhere else. Ride-sharing apps changed expectations for transportation, but they also changed expectations for food delivery, banking, healthcare, and government services. Once people became accustomed to tracking a driver in real time, waiting without visibility felt outdated in every industry. The same pattern is unfolding with intelligent software.

A person who spends an hour collaborating with ChatGPT, Claude, Gemini, or another intelligent assistant doesn't leave those expectations behind when opening a banking app, a CRM, an HR system, or an e-commerce website. Their expectations travel with them. They begin asking questions they never asked before. Why can't this application understand what I'm trying to do? Why do I have to fill in all these fields?

Why do I have to explain this again? Why can't it make a recommendation? Why does it forget everything after I close it? These questions are no longer about features. They are about expectations.

The Shift from Execution to Judgment


One of the most profound behavioural changes introduced by intelligent software is where people choose to spend their mental energy. Historically, software helped users execute work. Users gathered information, interpreted it, organized it, made decisions, and finally used software to record or execute those decisions. Increasingly, software performs many of those intermediate cognitive steps itself. It summarizes documents before we

read them. (See Fig 7.)

It drafts emails before we write them. It identifies patterns before we notice them. It proposes designs before we sketch them. It recommends code before we type it. People are spending less time generating possibilities and more time evaluating them. The role of the human is gradually shifting from execution toward judgment. That shift may prove to be one of the defining characteristics of the Intelligence Age.

Who does the cognitive work

FIG 7



The Blank Page Is Disappearing

For decades, software began with emptiness. A blank document. A blank presentation. An empty design canvas. An empty spreadsheet. The assumption was simple. The human would create. The software would assist. Thoughtware reverses that assumption. Increasingly, products begin with a thoughtful first draft. A pro-

posed itinerary. A suggested presentation structure. A generated interface. A summarized meeting. An initial implementation.

The blank page is becoming an exception rather than the default. This changes more than productivity. It changes confidence. People who once struggled to begin can now focus on improving. Professionals who once spent hours on repetitive drafting can spend those hours making higher-level decisions. The centre of gravity shifts from creation to refinement.

Learning Software Is No Longer the Goal

For decades, becoming proficient with software meant learning how it worked. People attended training sessions. Read manuals. Watched tutorials. Memorized shortcuts. Built mental models of where capabilities lived. Expert users were the ones who knew the interface best. Intelligent software changes that equation. The value increasingly lies not in learning the product, but in communicating the objective. Users care less about where a capability lives and more about whether the system understands what they want.

Learning software is gradually giving way to directing software. The interface becomes less of a destination and more of a conversation.

Consumers Expect Software to Remember

Memory is one of the clearest examples of changing expectations. Traditional software treated every interaction as largely independent. Each session began almost from scratch. Users repeatedly entered the same information, recreated the same preferences, and rebuilt the same context. They accepted this be-

cause software had no practical alternative. Thoughtware introduces a different expectation. Consumers increasingly assume that software remembers previous conversations, understands their preferences, recalls ongoing work, and builds upon what has already happened.

Forgetting no longer feels like a technical limitation. It feels like poor service. Just as people expect another human colleague to remember yesterday's discussion, they increasingly expect intelligent software to do the same. Memory is quietly becoming part of the user experience.

Consumers Expect Software to Understand Context

Memory alone is not enough. Remembering facts without understanding context is little more than storage. Consumers increasingly expect software to understand why they are asking a question, what they are trying to accomplish, what they have already attempted, and what constraints matter. When a user says, "Make it more professional," they rarely want to explain which document, which audience, or which previous discussion they are referring to.

Context has become implicit. Every time software asks users to repeat information they believe it should already know, trust declines. The most intelligent products often feel intelligent not because they answer better questions, but because they ask fewer unnecessary ones.

Consumers Are Becoming Comfortable Delegating

Perhaps the most surprising behavioural change is the willingness to delegate. People once hesitated to trust software with meaningful work. Today they routinely allow intelligent systems to draft important emails, summarize contracts, review code, recommend financial decisions, generate legal documents, plan trips, and organize research. This trust did not appear overnight. It developed gradually through thousands of successful interactions.

Consumers rarely begin by delegating important decisions. They begin by delegating small ones. As confidence grows, so does autonomy. This pattern is remarkably consistent with human relationships. Trust is earned incrementally.

The future of intelligent products will depend less on maximizing autonomy than on progressively earning it.


Consumers Judge Software Differently

These behavioural changes ultimately lead to a different set of evaluation criteria. Historically, consumers asked questions such as: Is it easy to learn? Is it fast? Is it reliable? Is the interface intuitive? Those questions still matter. But new ones are emerging. Does it understand me? Does it remember me? Can I trust its recommendations? Does it explain its reasoning? (See Fig 8.)

Does it reduce my thinking? Does it know when to ask and when to act? Increasingly, software is judged not only by usability, but by intelligence. That is a fundamental shift.

How consumers judge software

FIG 8



A New Consumer Demands New Products

Every technological revolution eventually produces a generation of consumers whose expectations are shaped by that technology from the very beginning. Children who grew up with smartphones never viewed mobile internet as a luxury. They viewed it as normal. The same will happen with intelligent software. The next generation will assume that products remember context. That software understands natural language.

That recommendations are personalized. That first drafts are generated automatically. That routine decisions are handled on

their behalf. To them, software that behaves otherwise will feel outdated. Just as command-line interfaces now feel foreign to most people, many of today's applications may eventually feel unnecessarily demanding. The products that succeed in the Intelligence Age will not merely incorporate artificial intelligence.

They will be designed around the psychology of consumers who have grown accustomed to collaborating with intelligence.

The Implications for Product Builders

This shift places an entirely new responsibility on those who build products. Designers can no longer think only about screens. Product managers can no longer think only about features. Engineers can no longer think only about implementation. Everyone involved in creating software must begin understanding how intelligent systems influence trust, memory, autonomy, reasoning, collaboration, and human decision-making. The profession itself is evolving.

Building products for the Intelligence Age requires a different mindset from the one that shaped the Information Age. The next chapter explores what that means for product leaders, designers, and engineers, and why understanding intelligence itself is becoming an essential part of building great products.

The Intelligence Product Leader

Every technological shift creates new products. The more interesting consequence is that it also creates new professions. The Industrial Revolution didn't simply introduce factories. It transformed the role of architects, engineers, managers, and craftsmen. The internet didn't merely create websites. It gave rise to entirely new disciplines such as interaction design, information architecture, search engine optimization, digital marketing, user experience research, and product management.

The Intelligence Age will do the same. The products we build are changing, but so are the people who build them. Product leaders, designers, researchers, and even engineers will need to develop a fundamentally different understanding of software. Not because their existing knowledge has become obsolete, but because another layer has been added to the profession. Just as every new age builds upon the previous one, every new generation of product builders builds upon the previous generation's skills.

The Foundations Still Matter

Whenever a new technology emerges, there is a temptation to believe everything that came before it no longer matters. History repeatedly proves otherwise. Good typography still matters. Visual hierarchy still matters. Accessibility still matters. Interaction design still matters. Information architecture still matters. User research still matters. Performance still matters. Empathy still matters. A product that is intelligent but confusing is still a poor product.

A product that reasons well but is visually overwhelming still creates friction. The fundamentals of great product design remain the foundation upon which intelligent experiences are built. The Intelligence Age does not replace traditional UX. It extends it.

Beyond Interface Design

For decades, the primary responsibility of designers was to shape interfaces. How should information be organized? Which button should users click? How many steps should a workflow contain? How should navigation behave? These questions remain important. But increasingly, they are no longer the most important questions. The Intelligence Age introduces a different class of design problems. What should the system remember?

When should it make a recommendation? How should it explain uncertainty? When should it ask for clarification? When should it take initiative? How much autonomy should it be given? What should it do when it is wrong? These questions are no longer about arranging pixels. They are about shaping intelligence.

Intelligence Becomes a Design Material

Designers have traditionally worked with materials such as typography, colour, spacing, motion, and layout. The Intelligence Age introduces entirely new design materials. Memory. Context. Reasoning. Planning. Recommendations. Autonomy. Conversation. Trust. These are not implementation details. They directly shape the experience users have with a product. A recommendation shown too early can be intrusive. A recommendation shown too late becomes useless.

Remembering too little feels forgetful. Remembering too much feels invasive. Explaining too much creates cognitive overload. Explaining too little reduces trust. Just as visual designers learn balance, hierarchy, rhythm, and proportion, future product designers will need to learn how to balance intelligence itself.

Understanding How Intelligence Works

Designing thoughtware requires a working understanding of the systems behind it. Not because every designer should become an AI engineer, but because good design depends upon understanding the capabilities and limitations of the medium being designed. Architects understand materials. Industrial designers understand manufacturing. Interaction designers understand software. Likewise, Intelligence Age product leaders must understand how intelligent systems behave.

That understanding begins with Large Language Models themselves. They need not understand the mathematics of transformers or attention mechanisms, but they should understand what language models are good at, where they struggle, how context windows influence behaviour, why hallucinations occur, how

latency affects experience, and how different models vary in capability, speed, and cost. Without this understanding, many product decisions become guesses.

Thinking Beyond the Model

Large Language Models are only one component of intelligent products. Most successful AI products are systems rather than models. The model reasons. Memory preserves knowledge across time. Retrieval provides current and relevant information. Tools extend capability beyond language. Agents coordinate complex workflows. Evaluation determines whether outputs are good enough. Designers and product managers increasingly need architectural literacy. They do not need to build these systems themselves. Each of those components is assembled from cognitive units, and each unit can be improved without rebuilding the product around it.

But they need to understand how they work together because each layer changes the user experience. A product that forgets everything after every session creates a different experience from one that builds long-term memory. A product grounded in retrieval behaves differently from one relying solely on model knowledge. A product with tools behaves differently from one limited to conversation.

These are product decisions before they are engineering decisions.

Designing for Memory

Memory deserves particular attention because it fundamentally changes the relationship between users and software. Traditional applications stored data. Thoughtware builds relationships. Future product teams will need to think carefully about what should be remembered, what should remain temporary, what should be forgotten, and how memory should evolve over time. They will need to distinguish between working memory, longterm memory, user preferences, project context, and organizational knowledge.

Memory is no longer simply a database feature. It is part of the user experience.

Designing Retrieval Instead of Content

For decades, product teams focused on organizing information. Increasingly, they will focus on retrieving the right information at the right moment. Semantic search. Embeddings. Retrieval-Augmented Generation. Hybrid search. Knowledge freshness. Grounding. Citation. These concepts may sound deeply technical today. In a few years, they will become as familiar to product teams as responsive layouts or information architecture are today.

The challenge is no longer storing information. It is delivering the right information precisely when it is needed.

Designing Agents

Perhaps the biggest conceptual shift lies in agentic systems. Traditional software waited. Users initiated every interaction. Agents behave differently. They pursue objectives. They create plans. They invoke tools. They recover from failures. They request help when necessary. They continue working until the objective has been achieved. Designing an agent is therefore less about designing screens and more about designing behaviour.

What decisions should the agent make independently? When should it seek approval? How should it communicate progress? How should it recover from mistakes? How should it gracefully fail? These questions increasingly define product quality.

A caution belongs here, because the word agent is applied far too widely. Much of what teams describe as an agent is a skill: a procedure whose shape somebody already knows, carrying real judgement inside it. That is not a lesser thing to have built. It is usually the more appropriate one, because a known procedure is easier to explain, cheaper to run and simpler to trust. The design question is not how to make a product more agentic. It is which parts of the work are settled enough to be performed, and which genuinely have to be worked out while the user waits.

Designing Trust

As intelligent systems assume greater responsibility, trust becomes one of the central challenges of product design. Users do not simply need correct answers. They need confidence that the system knows when it is uncertain. Future products will increasingly need mechanisms for explaining reasoning, citing sources, exposing confidence, requesting clarification, and enabling re-

covery without destroying progress. Trust is no longer created solely through visual polish. (See Fig 9.)

It is created through predictable intelligent behaviour. Consumers are remarkably forgiving of occasional mistakes. They are far less forgiving of overconfidence, inconsistency, or unexplained behaviour.

The intelligence product leader's competency map

FIG 9



Measuring Different Things


The metrics of successful products are also changing. Traditional product teams optimized completion rates, click-through rates, retention, engagement, and conversion. These metrics remain valuable. But intelligent products introduce entirely new questions. How often are recommendations accepted? How much cognitive effort has been eliminated? How often does the model re-

quire clarification? How accurate is memory? How frequently does retrieval improve responses? (See Fig 10.)

How often does autonomy succeed without intervention? How quickly does trust recover after failure? These are becoming design metrics as much as engineering metrics.

Activity metrics vs cognitive-outcome metrics

FIG 10




The Future Product Leader

The Intelligence Age is creating a new kind of product leader. Someone who is equally comfortable discussing customer psychology, interaction design, language models, memory architecture, retrieval systems, agent behaviour, evaluation strategies, and business outcomes. Not because they are expected to master every discipline, but because intelligent products blur the boundaries between them. Design decisions increasingly become architectural decisions.

Architectural decisions increasingly become experience decisions. The most successful product leaders will no longer think primarily in terms of screens or features. They will think in terms of intelligence. How should intelligence behave? Where should it help? Where should it wait? What should it remember? When should it recommend? When should it act? Those questions will increasingly define great products.

The four chapters in this part have established the foundation of this book. A new capability has emerged. Software has evolved to harness it. Consumers are changing because of it. And the people building products must evolve alongside it. The remainder of this book explores the philosophy that naturally follows. Not a philosophy of interfaces. A philosophy of designing intelligence.



PART II

The Philosophy

FIG 11

The twelve principles

Twelve principles. Four themes. Use this map to navigate Part II.
Each principle is simple to state, hard to master, and essential to remember.



The previous chapters argued that we are entering a new age. A new capability has emerged. Software has evolved to harness it.

Consumers are changing because of it. And the people building products must evolve alongside it. The natural question is: What should products built for this new age actually look like? Every technological age eventually develops its own philosophy of design.

The Industrial Age gave rise to industrial design because machines transformed how products were manufactured. The Information Age gave rise to interaction design and modern user experience because computers transformed how humans interacted with software. The Intelligence Age demands another evolution. Not because aesthetics have become less important. Not because usability no longer matters. But because intelligence has become a new design material.

These principles are not intended to replace traditional UX principles. A product with poor usability, weak accessibility, confusing navigation, or inconsistent interactions will not become delightful simply because it contains artificial intelligence. Instead, these principles build upon those foundations. They describe the new expectations emerging as consumers transition from operating software to directing intelligence. The principles are organized into four themes.

The first explores how intelligence should reduce cognitive effort. The second examines how products should build understanding over time through context and memory. The third focuses on intelligence itself as a user experience, exploring how software should reason, recommend, summarize, and explain. The fourth considers how interaction itself is changing, moving from operating interfaces toward collaborating with intelligent systems.

Together, these principles form a philosophy for designing products in the Intelligence Age.

Reducing Cognitive Load

For decades, digital products became increasingly powerful. Unfortunately, they also became increasingly demanding. Every new capability introduced another screen, another workflow, another menu, another setting, or another decision for users to make. The burden of complexity remained largely with the human. Good user experience design worked hard to reduce that burden. It simplified navigation. Reduced unnecessary clicks. Organised information.

Improved discoverability. Made interfaces more intuitive. Yet beneath all of those improvements remained one underlying assumption:

The software would provide tools. The user would provide cognition. That assumption is now changing. Intelligent software is capable of removing not only physical effort but cognitive effort. It can organize information before users see it, identify patterns be-

fore users notice them, recommend alternatives before users search for them, and prepare meaningful first drafts before users begin creating.

The first responsibility of intelligence is therefore not to impress users with sophisticated capabilities. It is to reduce unnecessary thinking. This theme explores four principles that fundamentally reduce cognitive effort while preserving human judgment.

Low-Cognitive First Draft

One of the most overlooked changes introduced by intelligent software is the disappearance of the blank page. For decades, almost every productivity tool began the same way. A blank document. A blank presentation. An empty spreadsheet. An empty design canvas. An empty email. The assumption behind every one of these products was simple: humans create, software assists. This assumption shaped generations of digital products.

Software provided a workspace. The user supplied the ideas. The software executed instructions. That model made perfect sense in the Information Age. It makes far less sense in the Intelligence Age.

The Blank Page Is Expensive

The blank page has never represented freedom. It represents work. Beginning a task often requires significantly more cognitive effort than continuing one. Before writing even the first sentence, people must organize their thoughts, decide where to begin, establish structure, overcome uncertainty, and commit to a direction. Psychologists sometimes refer to this as activation energy.


The first step is frequently the hardest. (See Fig 12.)

Once momentum exists, progress becomes dramatically easier. This is why people procrastinate before starting important work while finding it relatively easy to continue once they have begun. The blank page asks users to perform maximum thinking before any value has been created. Thoughtware changes that equation.

FIG 12

The cost of the blank page

A first draft removes the front-loaded cognitive effort so people can make progress, sooner.



From Creation to Refinement

One of the most significant behavioural shifts of the Intelligence Age is that people increasingly begin with something instead of nothing. A draft. An outline. A recommendation. A proposed itinerary. A generated interface. A suggested implementation. The human no longer starts by creating. The human starts by evaluating. This is a subtle change with profound implications. Evaluation is cognitively easier than creation.

Humans are remarkably good at recognizing what feels right, identifying weaknesses, spotting inconsistencies, and refining existing ideas. Beginning with a thoughtful proposal allows users to spend their mental energy on judgment instead of generation. The role of software shifts from providing a canvas to providing momentum.

Every Product Has a First Draft

Many teams interpret AI-generated first drafts too narrowly. They imagine writing assistants or code generators. The principle is much broader than that. A travel application can begin with a suggested itinerary instead of an empty planner. A CRM can propose likely follow-up actions instead of showing only customer records. A recruiting platform can generate interview questions before the interviewer begins typing.

A design application can propose multiple layouts instead of presenting an empty artboard. A healthcare system can prepare a draft clinical summary before the physician begins documentation. A project management tool can generate milestones and workstreams from a single objective. The implementation differs. The philosophy remains the same.

Whenever software can reasonably reduce the effort required to begin, it should.

Good First Drafts Are Not Final Answers


This principle is often misunderstood. The goal is not to replace human thinking. It is to elevate it. A poor first draft that demands complete rewriting provides little value. A thoughtful first draft creates momentum while still inviting human judgment. The draft should feel collaborative rather than authoritative. It should encourage improvement rather than discourage participation. Consumers increasingly expect software to meet them halfway. (See Fig 13.)

Not to finish the journey without them.

FIG 13

Calibrating the first draft

The goal is not to be perfect. The goal is to be useful.
Too vague fails. Too finished fails. The sweet spot enables progress.



Designing the First Draft

Designing effective first drafts requires restraint. The system should understand enough context to produce something relevant, but not so much that users feel they have lost ownership. The draft should expose assumptions instead of hiding them. It should be easy to modify, reject, regenerate, or redirect. Most importantly, it should reduce cognitive effort without reducing agency. The user should never feel trapped by the system's suggestion.

Instead, the first draft should feel like the opening move in a conversation.

The New Default

Throughout the Information Age, the blank page became one of the most familiar interfaces in software. The Intelligence Age quietly replaces that interface with something else. A thoughtful beginning. Consumers are already becoming accustomed to software that starts the work before they do. As this expectation spreads, products that continue presenting empty canvases for tasks that intelligence could reasonably initiate will begin to feel incomplete.

The blank page will not disappear entirely. There will always be moments that demand original thought from a human mind. But increasingly, those moments will become the exception rather than the rule. The first responsibility of intelligent software is not to finish the work. It is to remove the unnecessary effort required to begin.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Structured output	Generation constrained to a schema. The draft arrives as populated fields your UI can render, not prose you must parse.
Constrained decoding	Restricting the token space during generation. Guarantees valid shape on the first attempt rather than repair afterwards.
Prompt template	A parameterised, versioned prompt. The scaffold making drafts consistent across users and sessions.
In-context learning	The model adapting from examples in the prompt. Three or four good examples shape a draft further than a page of instruction.
RAG	Retrieving material and placing it in the prompt. The draft starts from the user's own documents rather than generic filler.
Prompt caching	Reusing computation for a stable prompt prefix. Drafting repeats the same long instructions, so caching cuts cost and first-token delay.
Time to first token	Delay before the first output appears. Governs whether the draft feels instant. The number that decides perceived speed.
Tokens per second	Sustained generation rate. Governs whether the draft finishes before attention drifts.
Streaming	Emitting tokens as they are produced. Turns a wait into visible progress.
Small language models	Compact models, far cheaper and faster. A draft is the cheapest place to use one, reserving the strong model for refinement.

Techniques

TECHNIQUE	WHAT IT MEANS
Model routing	Directing requests to different models by need. Draft on the fast model, refine on the capable one, within a single feature.
Speculative decoding	A small model proposes, a large one verifies. Raises tokens per second without changing which model the user effectively gets.
Sampling temperature	Randomness in token selection. Low for structured drafts, higher where the draft should offer a genuine alternative.
Seed from the user's material	Ground the draft in their most recent or relevant artefact before generating anything.
Ship it deliberately incomplete	Aim for relevant, structured, unfinished. Too polished and the user loses ownership.
Paint the skeleton first	Render structure immediately, fill sections as they stream.
Regenerate with a nudge	Offer a directional retry rather than a blank one.
Set a hard ceiling before first paint	Decide the maximum acceptable silence and engineer to it.
Measure effort saved	Track edit distance and discard rate, not accuracy.

Point & Fix

One of the oldest assumptions in software design is that changing something requires understanding how it was built. If users wanted to edit a document, they had to learn formatting tools. If they wanted to modify an image, they needed to understand layers, masks, and selection tools. If they wanted to change a website, they navigated complex editors. Even making a small adjustment often meant understanding the surrounding system.

The cognitive effort required to make a change was frequently disproportionate to the change itself. Consumers accepted this because software had no practical alternative. Intelligent software changes that assumption.

Increasingly, users expect to simply point at something and describe what should change.

The Cost of Explaining

Imagine reading a document and noticing a paragraph that feels too formal. Traditional software asks you to think like an editor. Select the text. Navigate the correct menu. Choose the appropriate command. Or manually rewrite it yourself. Now imagine the same interaction with intelligent software. Highlight the paragraph. “Make this friendlier.” Done. The difference is not merely speed. (See Fig 14.)

The difference is cognitive effort. The user never had to translate their intention into a sequence of software operations. They simply communicated the outcome they wanted. This pattern appears across almost every category of software. A designer circles a button and says, “Increase the spacing around this.” A developer selects a block of code and says, “Refactor this into reusable components.”

A marketer highlights a sentence and says, “Make this more persuasive.” An architect points to a room on a floor plan and says, “Move this wall two metres.” The interaction changes from operating software to directing intelligence.

FIG 14

Explaining vs pointing

The same edit expressed two ways.

Pointing carries the context. Explaining recreates it.

1 EXPLAINING (HIGH COGNITIVE COST)

Long global instruction that recreates what's already visible.

Q3 Product Update

Our active users grew 24% this quarter, driven by strong adoption across teams.

We shipped 18 improvements based on customer feedback.

Looking ahead, we're focused on **deepening integrations** and improving performance.

Change "deepening integrations" to "seamless integrations across tools" in the last paragraph.

High cognitive cost: You must describe what, where, and how—recreating context the system already sees.

2 POINTING (LOW COGNITIVE COST)

Selection provides the context. Five words express the intent.

Q3 Product Update

Our active users grew 24% this quarter, driven by strong adoption across teams.

We shipped 18 improvements based on customer feedback.

Looking ahead, we're focused on **deepening integrations** and improving performance.

seamless integrations across tools

WHAT THE SELECTION CARRIES IMPLICITLY

- Exact location (where)
- Surrounding context (what it relates to)
- Content type (phrase, not whole paragraph)
- Tone & style (match the rest)
- User's intent (improve, don't replace everything)

Low cognitive cost: You point to provide context. You only state the intent.

THE TAKEAWAY

Pointing lets users transfer context instead of retyping it. Less explanation. Less effort. Better results.

Precision Is Becoming the Default

One reason traditional software became increasingly complicated is that it often lacked context. The application rarely knew which part of the user's work required attention. Users therefore spent considerable effort identifying, navigating, and isolating the correct element before making a change. Modern intelligent systems understand context much more naturally. Selection itself becomes communication. A highlighted sentence. A selected image.

A marked region. A referenced screen. A voice instruction while pointing at an object. These all communicate context before a single word has been spoken. Pointing becomes part of the language. The user no longer explains where the problem exists. The act of pointing answers that question. The conversation can focus

entirely on what should change.

Local Changes Should Remain Local

One of the greatest strengths of point-and-fix interactions is that they naturally encourage surgical editing. When a designer asks an engineer to “change this button,” they do not expect the entire interface to be redesigned. When an editor says, “rewrite this paragraph,” they do not expect the rest of the document to change. Human collaboration already understands locality.

Thoughtware must learn the same discipline. One of the quickest ways to destroy trust is to modify parts of the work that the user never intended to touch. Consumers increasingly expect AI to make precise changes while preserving everything else. The untouched parts of the artifact are often just as important as the edited ones. Respecting that boundary is not merely a technical challenge.

It is a design principle.

Context Travels with the Selection

Selection does more than identify location. It communicates intent. Consider highlighting a sentence and asking, “Can you explain this?” The software immediately understands that “this” refers to the selected content. Highlight another paragraph and ask, “Rewrite this.” Circle an object in an image and say, “Remove this.” Select three files and ask, “Combine these.” The selection itself eliminates an entire layer of conversation.

Users no longer need to repeatedly explain which object they are referring to. The interface quietly provides the context. This dramatically reduces friction. It also makes intelligent interaction

feel surprisingly natural.

Beyond Text


Many people first encounter point-and-fix interactions through text editing. In reality, the principle is much broader. Design tools allow users to select visual elements and request changes. Development environments allow programmers to select code and request improvements. Presentation tools allow individual slides, charts, or images to be modified through conversation. Video editing increasingly enables users to point at specific frames or regions and describe desired changes. (See Fig 15.)

Three-dimensional design tools can allow designers to select objects and reshape them conversationally. Even the physical world is beginning to adopt the same interaction model through augmented reality and computer vision. The underlying principle remains unchanged. Users identify the object. Software determines how to modify it.

Point & fix across modalities

FIG 15

The power of pointing is that it works everywhere across the things users see, not just the things they write.



The End of Global Instructions

One reason prompt engineering often feels difficult is that many interactions remain unnecessarily global. Users ask software to improve an entire document when they really want one paragraph rewritten. They regenerate an entire image when only one object needs changing. They rebuild an interface because one component feels wrong. These are symptoms of software that lacks precision. As intelligent systems improve, consumers increasingly expect changes to remain local unless they explicitly request otherwise.

The scope of an instruction should match the scope of the user's intention. This is not simply an interface improvement. It reflects a deeper understanding of human communication. People rarely

describe every unaffected detail when requesting a change. They assume those details remain unchanged. Intelligent software should make the same assumption.

Designing for Point and Fix

Products built for the Intelligence Age should make pointing effortless. Selection should be lightweight. Referencing should feel natural. The user should never wonder whether the system understands what they are referring to. The boundary of change should be visible. The affected region should be obvious. The user should remain in control of accepting, rejecting, or refining the proposed modification.

Whenever possible, edits should feel reversible rather than destructive. The experience should resemble working with an experienced collaborator rather than issuing commands to a machine. A good collaborator understands what you are pointing at. A great collaborator also understands what you are not pointing at.

The New Editing Paradigm

The Information Age taught users to navigate software before making changes. The Intelligence Age teaches software to understand changes before requiring navigation. This may appear to be a small interaction improvement. It is not. It fundamentally changes how people approach editing. Instead of searching for the right tool, users simply identify the problem. Instead of describing entire workflows, they communicate desired outcomes.

Instead of rebuilding artifacts, they progressively refine them through small, precise conversations. Consumers are becoming

accustomed to software that understands exactly what they mean by “this.” As that expectation grows, products that continue requiring users to explain both the problem and its location will increasingly feel cumbersome. The future of editing is unlikely to revolve around learning more tools.

It will revolve around pointing at what matters and allowing intelligence to handle everything else.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Structured output	Schema-constrained generation. The model returns a patch (location, replacement, reason), not a rewritten document.
Constrained decoding	Restricting what may be generated. Enforces that only the selected span changes.
Context window	The working budget of the request. Send the selection and its neighbourhood, not the whole document.
Prompt caching	Reuse across calls sharing a prefix. Point-and-fix means many small edits against one document. Caching that document is the largest cost lever in this pattern.
Vision-language models	Models reading images alongside text. Lets a user point at a chart, a screenshot, or an interface element.
Segmentation models	Turning a click or box into a precise region. Converts a gesture into a region the model can reason about.
Document parsing	Layout-aware extraction of structure. Pointing inside a PDF or scan requires knowing where things sit on the page.
Logprobs	Token-level probabilities from the model. A usable signal for how sure the model is about an edit before it is applied.
Groundedness evaluation	Checking output against source. Here the check inverts: verify that what was not selected did not change.

Techniques

TECHNIQUE	WHAT IT MEANS
Return patches, not documents	Targeted changes are reviewable, revertible, cheap to verify.
Diff before applying	Machine-check that untouched content is untouched. Do not trust the instruction.
Send the neighbourhood	Include surrounding text, position, and structural role alongside the selection.
Anchor to stable identifiers	Character offsets break under concurrent edits. Node IDs survive.
Apply optimistically, undo cheaply	The edit lands instantly and reverses instantly.
Mark provenance	Track which changes came from the system so they can be reviewed as a set.

Intent over Instructions

One of the defining characteristics of software throughout the Information Age was that users had to tell it how to accomplish a task. Every interaction was a procedure. If you wanted to crop an image, you selected the crop tool, adjusted the boundaries, confirmed the action, and saved the result. If you wanted to sort a spreadsheet, you selected the range, chose the correct column, specified ascending or descending order, and applied the operation.

If you wanted to book a flight, you searched, filtered, compared, selected, entered passenger information, reviewed options, and finally completed the purchase. Software executed instructions. Humans defined the process. This arrangement made perfect sense because software had no practical way of understanding what people actually wanted to achieve. Its only option was to wait for explicit instructions. Thoughtware changes that relationship.

Increasingly, consumers communicate the outcome they want. Software determines how to achieve it.

Humans Naturally Think in Outcomes

Imagine asking a colleague to organize a meeting. You rarely describe every step. You don't say: (See Fig 16.)

“Open the calendar. Look for available slots. Compare everyone’s schedules. Reserve a meeting room. Create an invitation. Add the agenda. Send reminders.”


Instead, you simply say:

“Can you set up a meeting with the design team sometime next week?”

The colleague determines the process. Humans naturally communicate in terms of intent. For decades, software required us to communicate in terms of procedure. That mismatch has shaped nearly every digital interface we have ever used.

Where the translation work sits

Humans naturally think in outcomes



Translating Intent Was Always Work

Every interface asks users to perform an invisible task before performing the visible one. The invisible task is translation. The user begins with an objective.

“I want to compare these proposals.”

“I want to prepare for tomorrow’s interview.”

“I want to understand why revenue dropped.”

The software, however, understands only operations. Open. Import. Select. Filter. Sort. Configure. Export. The user therefore spends considerable effort translating intentions into operations the software can execute. The software rarely solves the original problem. It simply performs the requested operation. This translation layer is one of the largest sources of unnecessary cognitive

effort in modern software. Intelligence allows us to remove it.

Intelligence Owns the Procedure

Large Language Models introduce an entirely different interaction model. Instead of requiring users to define every step, intelligent systems increasingly infer an appropriate approach from the stated objective. Ask an intelligent assistant to prepare a project plan. It decides whether to begin with milestones, risks, stakeholders, timelines, or dependencies. Ask it to compare two contracts. It determines which clauses deserve attention.

Ask it to create a travel itinerary. It chooses an appropriate structure before generating recommendations. The user provides the destination. The system determines the route. This is one of the most important transfers of responsibility occurring in software today. Procedural thinking gradually moves from the human to the machine.

Intent Is Often Fuzzy

One concern frequently raised about intent-based systems is ambiguity. What if users don't know exactly what they want? The answer is surprisingly simple. Neither do humans. Consider how people collaborate. Requests are often incomplete.

"Make this feel more premium."

"I think something is missing here."

"Can we make this easier to understand?"

"This doesn't quite feel right."

These instructions are imprecise. Yet experienced colleagues rarely reject them because they lack specificity. They interpret

them. They ask clarifying questions when necessary. They propose alternatives. They gradually converge on the intended outcome. Intelligent software should behave similarly. The goal is not to eliminate ambiguity. It is to collaborate through ambiguity.

Instructions Become an Implementation Detail

Historically, users needed to understand how software worked because software did not understand them. As intelligent systems improve, procedural knowledge becomes increasingly unnecessary. A designer should not need to understand every manipulation required to achieve visual balance. A marketer should not need to know which menu generates audience segments. A manager should not need to know how scheduling algorithms allocate resources.

Consumers increasingly expect software to determine the implementation while they remain focused on the objective. This does not make expertise less valuable. It changes where expertise is applied. Experts increasingly judge solutions rather than manually constructing every step.

Designing Around Objectives

This shift has profound implications for product design. Traditional products often begin by exposing functionality. What actions can users perform? What menus are required? How should workflows be organized? Intent-driven products begin with a different question. What objectives bring people here? Someone opening a recruiting platform rarely wants to “create an interview template.” They want to hire the right person.

Someone opening a finance application rarely wants to “generate a report.” They want to understand what is happening in their business. Someone opening a healthcare system rarely wants to “complete documentation.” They want to treat a patient.

The closer products align themselves with human objectives, the less cognitive translation users must perform.

The Role of Clarification

Understanding intent does not mean pretending certainty. One of the strengths of human collaboration is knowing when clarification is necessary. If someone asks, (See Fig 17.)

“Book my flight to New York.”

A reasonable colleague immediately asks,

“Which New York?”

“When?”

“From where?”

“Do you have airline preferences?”


Clarification is not evidence of failure. It is evidence that the system understands the objective well enough to recognize missing information. Future products will increasingly differentiate themselves not by eliminating questions altogether, but by asking only the questions that genuinely matter. Every unnecessary question represents intelligence that failed to infer. Every necessary ques-

tion represents intelligence behaving responsibly.

When to ask

FIG 17

Ask only when it meaningfully improves the outcome. Don't interrogate users.



Intent Is Universal

Perhaps the most remarkable aspect of intent is that it transcends interfaces. The same objective can be expressed through typing, speaking, selecting text, uploading an image, pointing at an interface, or demonstrating an example. The interaction method changes. The intent does not. This suggests that future products should increasingly organize themselves around understanding objectives rather than collecting inputs.

Users should feel free to communicate however feels most natural. The responsibility for interpretation belongs to the software.

Designing for Intent

Designing intent-driven products requires a different mindset. The objective is no longer exposing every capability. It is understanding every objective. Products should identify the most common goals users arrive with and reduce the distance between intention and outcome. Interfaces should help software gather just enough context to act confidently. Clarification should appear only when truly necessary. Recommendations should reduce decision-making rather than increase it.

Most importantly, the product should remain focused on solving the user's problem rather than showcasing its own functionality. The best intelligent products will gradually make their internal complexity invisible. Consumers will increasingly judge them by how quickly they understand what needs to be accomplished. Not by how many features they contain.

From Features to Outcomes

The Information Age taught users to think in terms of software capabilities. The Intelligence Age encourages them to think in terms of outcomes. This shift may ultimately become one of the defining characteristics of intelligent products. Consumers are becoming less interested in learning what software can do. They are becoming more interested in expressing what they want to accomplish.

That changes the role of the interface. Instead of serving primarily as a catalogue of features, it becomes a medium for expressing intent. Once that intent is understood, software can determine the procedure. The user remains responsible for defining the destination. Intelligence increasingly takes responsibility for navigating

the path.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Function calling	The model emitting typed calls against defined tools. Turns a stated outcome into real operations in your product.
Tool schemas	Typed declarations of available capability. The model can only infer its way to actions you have described. Schema quality is product design.
MCP	An open protocol connecting models to tools and data. Standard access to capability, so intent can reach systems you did not build.
Agent loop	Plan, act, observe, revise, repeat. Outcomes usually need several steps, and a single call rarely satisfies an intent.
Task decomposition	Breaking a goal into ordered sub-goals. The translation work the user used to do, now done by the system.
Intent classification	Routing an utterance to a task type. Cheaper and more reliable than one large prompt attempting everything.
Structured output	Schema-constrained generation. Intent resolves to typed parameters, not a paraphrase.
Slot filling	Inferring missing parameters from available context. Determines how often you must ask versus assume.
Guardrails	Policy enforced outside the model. Bounds what the system may infer its way into doing.
Sandboxing	Executing against a safe replica. Lets an inferred plan be rehearsed before it touches anything real.

Techniques

TECHNIQUE	WHAT IT MEANS
Trajectory evaluation	Grading the path an agent took, not just its text. The only meaningful way to test whether intent was served.
Route before you reason	Classify the intent, then choose the handling path.
Infer defaults aggressively	Ask only for what genuinely cannot be derived from context.
Preview before executing	The cheapest place to catch a misread intent is before it acts.
One clarifying question, maximum	Clarification is an exception path. Interrogation is a failure.
Grade outcomes, not wording	Success is the result the user wanted, however the system got there.

Conversation over Configuration

Throughout the Information Age, software became increasingly configurable. As products grew more capable, they accumulated options. Preferences, settings, filters, toggles, permissions, rules, workflows, templates, dashboards, shortcuts, and automation builders gave users remarkable control over how software behaved. The intention was good. Different users had different needs, and configuration made software flexible enough to accommodate them. The cost, however, was complexity.

Every new capability required another decision. Every decision required another interface. Every interface required users to learn another part of the product. Over time, many applications became less limited by their capabilities than by their ability to expose those capabilities. Intelligent software changes that equation. Instead of asking users to configure behaviour in advance, it increasingly allows them to describe behaviour when they need it. Configuration gives way to conversation.

Configuration Is Deferred Thinking

Configuration has always required users to make decisions before they possess enough context to make them well. Consider setting up a project management tool. Before managing a single project, users are often asked to configure workflows, notification preferences, permissions, templates, labels, automations, dashboards, integrations, and dozens of other settings. Many of these decisions are guesses. Users cannot know what will work best until they have actually begun using the product. (See Fig 18.)

Software has traditionally solved this uncertainty by asking users to decide anyway. Conversation offers a different approach. Instead of predicting future needs, users simply describe what they need when the moment arrives. The system adapts accordingly.

FIG 18

Configuration predicts, conversation accumulates

Same outcome, different burden.

CONFIGURATION

All decisions demanded up front.

You must decide everything before you even know what you need.

Start
(blank slate)

FRONT-LOADED DECISIONS

Every option, setting, and rule must be chosen upfront.



Finally,
it's ready.




Settled state
The outcome you wanted.

CONVERSATION

Decisions made as they arise.

You decide only what's needed, when it's needed.

Start
(blank slate)



Naturally converges.



CONFIGURATION BURDEN

High upfront cognitive load.
Risk of wrong or unnecessary choices.



CONVERSATION BENEFIT

Low upfront burden.
Decisions accumulate into the right outcome.

Human Relationships Rarely Begin with Configuration

Imagine hiring a new colleague. You don't spend the first day defining every possible scenario they may encounter over the next five years. You don't write rules for how they should respond to every email, schedule every meeting, organize every document, or communicate every decision. Instead, you begin working together. You provide context. You answer questions. You make adjustments.

Over time, the collaboration naturally becomes more personalized. Software has historically followed the opposite model. Before meaningful work begins, users are expected to configure everything. Intelligent software increasingly behaves more like a colleague. It learns through interaction rather than demanding

complete specification upfront.

Conversation Is Progressive Configuration

Conversation does not eliminate configuration. It distributes it across time. This distinction is important. Products will always require some explicit preferences. Language. Time zone. Privacy settings. Notification permissions. Security policies. These decisions are stable and deserve explicit control. The problem arises when products ask users to configure dynamic behaviour that could instead emerge naturally through interaction. A user should not need to build a complicated automation simply to say,

“Whenever I receive invoices from this supplier, categorize them under operations.”

Nor should they need to navigate multiple configuration panels to express,

“Summarize every meeting longer than thirty minutes and send the summary to the participants.”

These are conversations. The product can convert them into persistent behaviour behind the scenes.

*Conversation becomes configuration
expressed in natural language.*

Configuration Becomes Adaptive

Traditional configuration assumes behaviour remains fixed until manually changed. Intelligent systems introduce a more adaptive model. Instead of asking users to continuously maintain rules, software can observe patterns and make suggestions. It might notice that every Monday morning a manager reviews sales performance and ask whether this summary should be prepared automatically. It might observe that a designer consistently prefers concise AI responses and gradually adjust future interactions.

It might recognize recurring project structures and suggest applying them to similar initiatives. The software evolves alongside the user instead of remaining frozen in the state in which it was originally configured. This creates a very different relationship between users and products.

The Best Settings Are Invisible

Many of today's most delightful digital experiences already hide enormous amounts of configuration. Smartphone cameras automatically adjust exposure, white balance, focus, image stabilization, and computational photography without asking users to configure each parameter. Navigation applications determine routes without exposing every optimization algorithm. Search engines rarely ask users how relevance should be calculated. The success of these products comes from making intelligent decisions on behalf of the user while still allowing manual control when necessary.

The same principle applies to intelligent software. Users increasingly care less about how behaviour is configured and more about

whether the behaviour aligns with their intent. Configuration becomes an implementation detail.

Conversation Builds Better Mental Models

Another advantage of conversation is that it aligns with how humans naturally think. People rarely think in menus. They think in goals.

“Help me prepare for this meeting.”

“Organize these files.”

“Find the differences between these contracts.”

“Make this presentation more persuasive.”

Conversation allows users to remain focused on the problem they are trying to solve. The software assumes responsibility for translating that goal into actions. This dramatically reduces the mental effort required to use sophisticated products. Instead of learning the product’s language, users communicate in their own.

When Configuration Still Matters

Conversation should not become an excuse to eliminate control. Some decisions deserve explicit visibility because they affect privacy, security, compliance, governance, or organizational policy. A financial approval workflow should not silently change itself. A healthcare system should not infer regulatory requirements through conversation alone. Enterprise software will continue requiring explicit configuration where predictability and accountability matter. The principle is not to eliminate settings. (See Fig 19.)

It is to reserve settings for decisions that genuinely benefit from permanence, visibility, and deliberate control. Everything else should increasingly become conversational.

What stays a setting

FIG 19

When configuration still matters.

	 KEEP AS EXPLICIT SETTINGS Configuration still matters here.	 MOVE TO CONVERSATION Make the default smart. Refine in dialogue.
 WHO & WHAT WE MUST PROTECT (High consequence)	Privacy, security, compliance, governance and organizational policy. Risks are too high and the rules don't change per conversation.	Everything else about doing the work. The system can ask, assume, and adapt as context unfolds.
 BENEFITS FROM PERMANENCE (Endures over time)	Decisions that should persist. Set once, remembered everywhere. Changes are infrequent and deliberate.	Decisions that are situational. Context changes. Preferences vary. Refine as you go.
 NEEDS VISIBILITY AND AUDITABILITY (Must be visible)	Decisions that need to be visible, reviewed and auditable. Others must rely on them and trust them. They form part of the record.	Working choices that don't need global visibility. Kept in context. Not worth adding to the settings surface.
 REQUIRES DELIBERATE CONTROL (Not accidental)	Decisions that have real-world impact and deserve deliberate control. Set intentionally, with clear ownership. Not something the system should guess.	Choices the system can intelligently propose, test and adjust. Keep the user in flow. Make change easy and reversible.
 PRINCIPLE Keep settings for what must be stable, visible, and controlled. Make everything else conversational.		OUTCOME Fewer knobs. Lower burden. Higher trust. The right decisions in the right place.

Designing Conversational Products

Designing conversational products requires restraint. Not every interaction should become a chat. Conversation is valuable when users are expressing intent, resolving ambiguity, or defining behaviour. Simple actions should remain simple. Pressing “Play” should still play a video. Checking a checkbox should still check a box. The goal is not to replace every interface with a conversation. The goal is to remove unnecessary configuration where conversation is a more natural way of expressing intent.

Good products will seamlessly combine both interaction models. Deterministic interfaces for deterministic actions. Conversation for ambiguous, flexible, or evolving ones.

From Configuration to Collaboration

The Information Age required users to configure software before software could help them. The Intelligence Age increasingly allows software to learn through collaboration. Consumers are becoming less willing to invest hours configuring products they have not yet used. Instead, they expect products to become more useful as the relationship develops. Conversation is not simply a replacement for settings.

It is a different philosophy of interaction. Rather than forcing users to predict every future need, it allows products to evolve alongside them. Configuration asks users to understand software. Conversation allows software to understand users. As intelligent systems become more capable, the most successful products will not be the ones with the largest settings pages. They will be the ones that require users to think about settings the least, while still giving them confidence that the product behaves exactly as they intend.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Structured output	Schema-constrained generation. An utterance becomes a config object, so conversation produces inspectable state.
Function calling	Typed calls the model can make. Settings become tools the system can set on the user's behalf.
Procedural memory	Stored knowledge of how the user prefers things done. Where a conversational preference actually lives once accepted.
Preference store	A structured, queryable record of user choices. Distinct from conversational memory: this part must be visible and editable.
Session vs persistent state	What lasts a turn, a task, or forever. Getting scope wrong is the most common failure in conversational configuration.
Memory conflict resolution	Reconciling instructions contradicting earlier ones. "Actually, always do X instead" must cleanly supersede, and visibly so.
Prompt template	Parameterised prompt construction. Accepted preferences get injected into every subsequent request.
Deterministic execution	Running settled rules as code, not inference. Once a preference is fixed, stop asking the model. Rules should be reliable, not re-reasoned.
Audit logging	An immutable record of state changes. For governance-relevant settings: who set what, when, via which utterance.

Techniques

TECHNIQUE	WHAT IT MEANS
Promote repetition into rules	When a one-off recurs, offer to make it standing.
Confirm before persisting	The moment a request becomes configuration should be explicit.
Always render created state	Conversation that writes invisible settings is worse than a settings panel.
Make supersession visible	Show what changed when a new instruction overrides an old one.
Keep governance out of dialogue	Privacy, security, compliance and policy stay as explicit controls.

Software That Knows You

The Information Age taught us to design software that responds correctly. The Intelligence Age demands software that understands. For decades, most digital products treated every interaction as an isolated event. A user opened an application, completed a task, closed it, and returned later to begin another task. Although software stored enormous amounts of data, very little of that data contributed to making future interactions meaningfully better.

The relationship remained largely transactional. Every session began almost from scratch. Users repeatedly entered the same information, recreated the same context, re-established the same preferences, and explained the same objectives. This behaviour became so common that people stopped questioning it. We accepted that software forgot. Human relationships work differently. Every meaningful relationship builds upon previous interactions. Friends remember conversations. Colleagues remember projects.

Doctors remember patients. Teachers remember students. The quality of the relationship improves because memory accumulates over time.

Consumers are beginning to expect software to behave in a similar way. They increasingly assume that products remember previous conversations, understand ongoing work, recognize preferences, and build upon knowledge instead of repeatedly discarding it. This represents one of the most important shifts in product design. Software is no longer judged solely by how well it responds in the moment.

It is increasingly judged by how well it understands the relationship over time. The next two principles explore the foundations of that relationship. The first concerns context, the ability to understand the situation surrounding a request. The second concerns memory, the ability to preserve understanding across time. Together, they transform software from a collection of interactions into an evolving partnership.

Context is a Feature

Traditional software was designed around inputs. Thoughtware must be designed around context. That distinction appears subtle. In practice, it changes almost every interaction users have with a product. For decades, software has largely treated context as supporting information. It existed to help complete a task but rarely influenced the product's behaviour in meaningful ways. Users selected a file, opened a document, filled in a form, or entered a search query. The application processed exactly what was presented to it.

Everything outside those explicit inputs was largely invisible. Intelligent software operates differently. Its quality depends not only on the request itself, but on everything surrounding that request.

Understanding context has become as important as understanding language.

Information Without Context Is Ambiguous

Imagine receiving a message that simply says,

“Let’s move it.”

Without context, the sentence is almost meaningless. What should be moved? Who is involved? When should it happen? Why is it being moved? Every human conversation depends on shared context. People rarely explain everything from the beginning because both participants already understand much of the surrounding situation. Software has traditionally lacked this ability. Consequently, users learned to over-explain.

They copied background information into every prompt. Repeated previous decisions. Uploaded the same documents. Re-stated the same objectives. The software wasn’t unintelligent. It was context-blind.

Consumers No Longer Want to Repeat Themselves

One of the clearest behavioural shifts of the Intelligence Age is the declining tolerance for repetition. Consumers increasingly expect software to know what they are currently working on. If they ask,

“Can you make this more concise?”

they expect the system to know what this refers to. If they ask,

“Continue where we left off,”

they assume previous work is still available. If they upload a design file and later ask,

“What do you think about the navigation?”

they do not expect to explain which design they mean. Repeating context feels increasingly unnecessary because consumers have experienced products capable of carrying conversations forward. Once people become accustomed to contextual understanding, its absence becomes frustrating.

Context Exists Everywhere


Many teams think of context only as conversation history. In reality, context exists at every level of a product. The user's current task. Their previous interactions. The document currently open. The selected text. The project they belong to. Their role within an organization. The device they are using. The location from which they are working. The time of day. (See Fig 20.)

The people collaborating with them. The goals they have previously expressed. The knowledge they have already demonstrated. Every one of these signals helps software understand what users actually mean rather than simply what they type. Great intelligent products continuously gather and organize these signals without forcing users to provide them repeatedly.

FIG 20

The context inventory

Context exists everywhere.



Context Reduces Questions

A useful way to think about intelligence is this: The more context a system understands, the fewer unnecessary questions it asks. Imagine asking an experienced colleague,

“Can you send that to the client?”

If both of you have been discussing the same proposal for the past hour, almost no clarification is required. Now imagine asking a stranger. The number of questions immediately increases. Which client? Which document? Which version? Should I send it today? By email? With any additional notes? Nothing about the request changed. Only the available context changed. The same principle applies to software.

Every unnecessary clarification represents missing context. Every avoided clarification represents successful understanding.

Context Should Be Gathered, Not Requested

Traditional software often waits for users to provide context explicitly. Thoughtware should increasingly gather it automatically. If the user has highlighted text, that selection becomes context. If they are editing a presentation, the current slide becomes context. If they are reviewing code, the repository becomes context. If they are discussing a customer account, the CRM record becomes context.

The product should quietly assemble the surrounding picture before asking users to explain it again. The responsibility for gathering context gradually shifts from the human to the software. This is one of the defining characteristics of intelligent products.

Context Is Temporary

One important distinction must be made. Context is not memory. Context answers the question, (See Fig 3.1)

“What is happening right now?”

Memory answers,

“What should still matter tomorrow?”


A user editing a presentation may want the system to understand the previous slides, the audience, and the objective of today’s meeting. That information is highly relevant now. It may become irrelevant once the presentation is complete. Context is often temporary. It expands and contracts as work progresses. Understanding this distinction prevents products from remembering everything unnecessarily while still remaining deeply aware of the current situation.

The next chapter explores how to determine what should persist beyond the present moment.

Context half-life

FIG 3.1

All context decays. The longer it goes unused, the less reliable it becomes.
Fresh context keeps intelligence grounded. Expired context creates risk.



Context Is a Product Capability

Historically, context has often been treated as an implementation detail. Something engineers assembled behind the scenes. The Intelligence Age changes that perspective. Context directly shapes the quality of the experience. Whether software understands a user's objective. Whether it asks unnecessary questions. Whether it retrieves relevant information. Whether it makes appropriate recommendations. Whether it acts confidently. These are product outcomes.

Consequently, context can no longer be viewed as a technical optimization. It has become a first-class product capability. Design-

ers should ask where context comes from. Product managers should determine which contextual signals matter. Engineers should build systems capable of maintaining it. Researchers should understand which context consumers naturally expect software to possess. Every discipline now contributes to contextual intelligence.

Designing for Context

Products designed for the Intelligence Age should assume that every unnecessary request for context represents friction. The objective is not to collect more information. It is to understand more from the information already available. Every interaction should begin with the question,

“What does the product already know that the user should not have to repeat?”

That simple question often reveals remarkable opportunities to improve an experience. Consumers rarely notice when context has been handled well. The interaction simply feels effortless. They notice immediately when it has not. The best contextual systems quietly disappear into the background. Users stop thinking about how the software understands them. They simply expect that it does. That is precisely when context has become a feature rather than merely supporting information.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Context window	Maximum tokens a model can attend to. The hard budget every other decision in this chapter spends against.
Context engineering	Deciding what enters the prompt, in what order. The central discipline. Not plumbing, but where product judgment lives.
Embedding models	Models mapping text to vectors capturing meaning. The substrate of retrieval. Choice of embedding model bounds retrieval quality.
Vector database	Storage and nearest-neighbour search over embeddings. Where retrievable context lives at scale.
Semantic search	Retrieval by meaning rather than wording. Finds relevant material the user did not name exactly.
Lexical search (BM25)	Classic keyword ranking. Still better than semantic search for names, codes, and exact identifiers.
Hybrid search	Combining lexical and semantic scoring. Reliably outperforms either alone. The pragmatic default.
Reranking	A cross-encoder reordering retrieved candidates. Usually the highest-leverage single improvement to a retrieval system.
Chunking strategy	How documents are split before indexing. Determines what can ever be found. No model recovers from bad chunking.
Metadata filtering	Constraining retrieval by structured attributes. Scopes context to the right project, period, or document type.

CONCEPT	WHAT IT MEANS
Knowledge graph	Entities and relationships as explicit structure. Answers questions about how things relate, which vector search handles poorly.
GraphRAG	Retrieval traversing a graph rather than a flat index. For multi-hop questions where the answer spans several connected documents.
Entity resolution	Recognising that references point to the same thing. Without it, context about one customer fragments across name variants.
RAG	Retrieval-augmented generation. The overall pattern this chapter describes in prose.
Agentic RAG	The model deciding what to retrieve, iteratively. Handles questions where the right query is not known upfront.
MCP	Standard protocol for tool and data connection. How context reaches the model from systems you do not own.
Permission-aware retrieval	Filtering by entitlement at query time. Filtering after retrieval leaks. This must happen in the query.
Prompt caching	Reusing a stable prompt prefix. Makes large persistent context economically viable turn after turn.
Context degradation	Reduced attention to material mid-window. More context is not better. Placement affects what the model actually uses.
Grounding and citation	Tying output to retrieved spans. Turns retrieved context into something a user can verify.
Index freshness	How current the retrievable corpus is. Stale context is more dangerous than absent context.

Techniques

TECHNIQUE	WHAT IT MEANS
Retrieve wide, rerank hard	Fetch generously, then let a reranker decide what enters the window.
Filter permissions in the query	Never retrieve then redact.
Budget the window per turn	Decide explicitly what share goes to instructions, memory, retrieval, history.
Place decisive material at the edges	Beginning and end are attended to more reliably than the middle.
Measure retrieval separately	Evaluate whether the right context was found before blaming the model.

Memory is Part of UX

Context helps software understand the present. Memory helps software improve the future. Although the two are closely related, they solve fundamentally different problems. Context answers the question, “What is happening right now?” Memory answers a different one: “What should the software continue to know about me after today?” This distinction may appear technical. It is not. It is one of the most important user experience decisions intelligent products will make.

For decades, software has been remarkably forgetful. Every new project, every new conversation, every new session required users to rebuild understanding from scratch. We accepted this because software had little ability to accumulate knowledge in meaningful ways. Consumers are beginning to expect something very different. They expect software to grow with them.

Great Relationships Have Memory

Imagine meeting the same colleague every morning and having to introduce yourself again. You explain your role. The projects you're working on. The terminology your team uses. Your preferences. Your goals. The conversation ends. The next morning, the colleague remembers none of it. The relationship would never deepen. The problem would not be intelligence. It would be memory.

Human relationships become valuable because they accumulate understanding over time. Every conversation adds another layer of shared knowledge. Future interactions become easier because less time is spent rebuilding context. Thoughtware is beginning to follow the same pattern. Consumers increasingly expect every interaction to make the next interaction better.

Software Has Always Stored Data

One common misconception is that software has always had memory. Technically, that is true. Products have stored customer records, documents, preferences, and databases for decades. But storing data is not the same as remembering. A CRM stores thousands of customer interactions. A colleague remembers that a particular customer dislikes long meetings. A project management tool stores completed tasks.

A teammate remembers which style of communication works best with each stakeholder. Memory is not about retention. It is about useful understanding. Consumers increasingly expect software to remember in the same way people remember. Not every detail. Only the details that improve future collaboration.

Not Everything Should Be Remembered


Perhaps the most difficult question intelligent products face is not what they can remember. It is what they should remember. Remembering everything creates clutter. It also creates discomfort. Human memory is selective. We remember recurring preferences. Long-term goals. Important decisions. Relationships. Significant events. We forget countless temporary details because they no longer contribute to future interactions. Intelligent products should behave similarly. (See Fig 21.)

A temporary request to summarize a document probably should not become permanent memory. A long-term writing preference probably should. A travel itinerary for next week may become irrelevant once the trip ends. A preferred seat on flights may remain valuable for years. Memory is not accumulation. It is curation.

FIG 21

Remember, ask, or forget

Not everything should be remembered.



Memory Should Reduce Repetition

The primary purpose of memory is remarkably simple. It should eliminate unnecessary repetition. Users should not repeatedly explain how they like responses structured. They should not continually restate ongoing projects. They should not repeatedly identify their preferred tone, preferred language, favourite tools, recurring collaborators, or organizational context. Every time

software asks users to repeat something they reasonably expected it to remember, the relationship becomes weaker.

Memory reduces friction not by adding more information, but by removing unnecessary conversations.

Memory Should Be Visible

One reason people trust human memory is that they can challenge it. If a colleague says,

“I remember you preferred the earlier proposal,”

you can immediately respond,

“Actually, I changed my mind.”

Software should offer the same opportunity. Memory should never feel mysterious. Users should understand what the system remembers, why it remembers it, and how they can change or remove it. Invisible memory often feels intrusive. Visible memory feels collaborative. People are surprisingly comfortable with products remembering them when they remain in control of that relationship.

*Transparency transforms memory from
surveillance into service.*

Memory Evolves


People change. Preferences change. Goals change. Projects end. Relationships evolve. Memory that never changes eventually becomes inaccurate. One of the characteristics of good human memory is adaptation. We naturally update our understanding as we learn more about someone. Intelligent products must do the same. Memory should be continuously refined rather than permanently fixed. Products should recognize when new behaviour consistently contradicts old assumptions.

They should ask for confirmation before updating important memories. They should gracefully forget information that no longer appears relevant. A living memory is more valuable than a permanent one.

Memory Creates Compounding Experiences

Perhaps the most remarkable characteristic of memory is that its value compounds over time. The first interaction with an intelligent product may feel only slightly better than traditional software. The hundredth interaction should feel dramatically different. The software understands terminology without explanation. It anticipates preferred formats. It recognizes recurring objectives. It recalls previous decisions. It recommends based on accumulated experience rather than isolated requests. (See Fig 22.)

The relationship becomes progressively more valuable because every interaction contributes to the next. Traditional software delivers roughly the same experience every day. Intelligent products should become better collaborators over time.



Designing Memory

Designing memory requires answering questions that product teams rarely asked in the Information Age. What deserves longterm memory? What should remain temporary? Who owns that memory? How is it updated? How can users inspect it? How can they edit it? How can they delete it? Should memory belong to an individual, a team, an organization, or a project?

Should it travel across devices? Across applications? Across conversations? These are not implementation details. They shape the relationship between consumers and products for years. Future product teams will spend as much time designing memory as previous generations spent designing navigation.

Memory Is Part of User Experience

For decades, user experience focused primarily on what happened during an interaction. The Intelligence Age extends that responsibility across interactions. The quality of today's experience increasingly depends upon what happened yesterday. A product that remembers appropriately feels thoughtful. One that forgets important information feels indifferent. One that remembers the wrong things feels intrusive. Designing memory therefore becomes an exercise in designing relationships.

Consumers no longer judge products solely by how well they perform a task today. They increasingly judge them by whether tomorrow's experience is better because of today's interaction. That is why memory is no longer simply a technical capability. It is part of the user experience itself. And once products begin understanding the present through context and improving the future through memory, they become capable of something even more valuable.

They become capable of helping users make better decisions. The next theme explores how intelligence itself becomes part of the experience, not merely through answers, but through recommendations, reasoning, and judgment.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Short-term memory	State held within a session. Cheap, bounded, discarded. Most products stop here.
Long-term memory	State persisting across sessions. What the chapter is actually arguing for.
Episodic memory	Records of specific past events and interactions. "You rejected this supplier last quarter".
Semantic memory	Durable facts about the user or their world. Role, team, standing constraints, stable attributes.
Procedural memory	How this user prefers work to be done. Tone, format, defaults, working style.
Memory extraction	Deriving durable facts from raw interaction. The pipeline turning transcripts into something worth keeping.
Salience and write policy	Deciding what merits storage at all. Most systems store far too much and recall far too little.
Summarization compaction	Condensing history to fit a budget. Keeps long relationships affordable, at the cost of detail.
Embedding models	Vector representations for similarity. How a relevant memory gets found when needed.
Vector store vs profile store	Fuzzy recall versus structured facts. Preferences belong in a structured, editable record. Episodes suit vector recall.
TTL and decay	Ageing and expiry rules for stored items. Memory without an ageing model becomes noise.
Conflict resolution	Handling new facts contradicting stored ones. Silently keeping both is the most common bug in memory systems.

Techniques

TECHNIQUE	WHAT IT MEANS
User-editable memory	Interfaces to view, correct, and delete. The entire difference between memory and surveillance.
Data residency and erasure	Where personal data lives and how it is removed. A product constraint on memory design, not a legal footnote.
Write less, recall better	Storage is easy. Surfacing the right memory at the right moment is the hard problem.
Age deliberately	Give every memory type an expected lifespan when you design it.
Separate stable from volatile	A seat preference and next week's itinerary need different treatment.
Show the memory	Make what is remembered visible and correctable in the product.
Ask before persisting anything sensitive	Consent at the moment of capture, not buried in a policy.

Intelligence as the Interface

Throughout the Information Age, the interface was the product. Design teams invested enormous effort into arranging information, simplifying navigation, reducing clicks, and exposing functionality. Success depended largely on how efficiently users could move through an interface to accomplish their goals. The interface was where value was delivered. The Intelligence Age introduces a profound shift. The interface is no longer the primary experience.

Intelligence is. Consumers are increasingly less concerned with how they interact with software and more concerned with how well the software thinks alongside them. A beautifully designed inter-

face that offers poor recommendations, forgets important context, or requires users to think through every decision will often feel less valuable than a simpler interface that consistently helps

users make better choices.

This does not diminish the importance of interface design. Rather, it changes where differentiation comes from. Visual design is becoming a commodity. Intelligence is becoming the experience. That experience is expressed through a number of behaviours. How well the product recommends instead of merely responds. How effectively it reduces cognitive effort. How transparently it explains its reasoning. How intelligently it applies its capabilities.

The next four principles explore these new dimensions of product experience.

Recommend, Don't Respond

One of the greatest limitations of traditional software is that it waits. It waits for the user to know what they want. It waits for the user to find the right feature. It waits for the user to ask the right question. It waits for the user to notice the problem. Only then does it respond. This model made sense because software had little ability to understand situations beyond the explicit commands it received.

Intelligent software changes that relationship. Instead of waiting to be instructed, it can increasingly recognize opportunities to help before users ask. The role of software begins shifting from responding to recommending.

Consumers Don't Always Know What to Ask

Most people approach software with incomplete information. A first-time founder opening a financial dashboard may not know which metrics deserve attention. A doctor reviewing patient records may not immediately notice subtle trends. A designer may sense that a layout feels unbalanced without knowing exactly why. A hiring manager may struggle to compare candidates fairly across dozens of interviews.

The limitation is not intelligence. It is perspective. People often know their goal without knowing the best path toward it. Good colleagues recognize this. They don't simply answer questions. They volunteer observations.

"Have you considered this?"

"This number seems unusual."

"You may want to review this section first."

"I think you're optimizing the wrong thing."


Recommendations emerge because understanding extends beyond the question itself. Software should increasingly behave the same way.

Answering Questions Is the Lowest Level of Intelligence

Search engines answered questions. Large Language Models answer questions more naturally. Neither represents the highest expression of intelligence. The real value emerges when software helps users ask better questions, or removes the need to ask them altogether. Consider a meeting assistant. A traditional assistant

transcribes the meeting because that is what the user requested. An intelligent assistant notices unresolved decisions, identifies conflicting opinions, proposes follow-up actions, highlights missing stakeholders, and reminds participants about commitments made during previous meetings. (See Fig 23.)

The transcript becomes secondary. The recommendation becomes valuable. Intelligence is measured less by how accurately software responds and more by how effectively it improves decision-making.



Recommendations Reduce Cognitive Blind Spots

Every decision contains blind spots. Humans naturally focus on information they already consider important. They overlook patterns. Ignore contradictions. Miss opportunities. Forget previous experiences. Recommendations compensate for these limitations. A product reviewing a contract might recommend examining an unusually broad liability clause. A writing assistant might recommend simplifying an argument that has become unnecessarily complex. A recruiting platform might recommend additional interview questions because a candidate's experience appears inconsistent.

A project management system might recommend revisiting timelines after detecting unrealistic dependencies. In each case, the software contributes perspective rather than simply execution. It helps users see what they may not have noticed themselves.

Recommendation Is Not Automation

One of the most common misconceptions about intelligent products is that recommendations inevitably lead toward full autonomy. The two are fundamentally different. Automation removes humans from decisions. Recommendations improve human decisions. A recommendation preserves agency. The user remains responsible for accepting, rejecting, modifying, or ignoring the suggestion. This distinction is essential. Consumers are often remarkably comfortable receiving recommendations long before they are comfortable delegating decisions.

Recommendation therefore becomes an important intermediate stage in building trust. It demonstrates competence without re-

moving control.

Great Recommendations Arrive at the Right Moment

A recommendation is valuable only if it appears at the appropriate time. Recommendations delivered too early interrupt thinking. Recommendations delivered too late become obvious. The timing of intelligence is therefore just as important as its quality. Consider navigation software. Suggesting an alternate route before heavy traffic forms is useful. Suggesting it after the driver has already entered the congestion is considerably less valuable.

The same principle applies to every intelligent product. Recommendations should appear when they reduce effort rather than increase interruption. Designing recommendation timing becomes as important as designing recommendation quality.

Recommendations Must Respect Confidence

Not every observation deserves to become a recommendation. Thoughtware should understand the difference between confidence and speculation. When confidence is high, recommendations can be presented proactively. When confidence is lower, software may instead ask clarifying questions or present alternative possibilities. Consumers are surprisingly tolerant of uncertainty when it is communicated honestly. What they dislike is false confidence. (See Fig 24.)


Products should therefore recommend only when the expected benefit outweighs the cost of interruption.

Silence is often better than poor advice.

Confidence determines behaviour

FIG 24

Recommendations must respect confidence



Designing Recommendations

Recommendation is not merely an AI capability. It is a product capability. Product teams must decide what deserves recommendation, when recommendations should appear, how strongly they should be expressed, and how users should interact with them. Should recommendations interrupt the user? Should they remain passive until requested? Should they explain their reasoning? Should they adapt based on previous acceptance or rejection?

Should they disappear once ignored? These are experience decisions. A recommendation that consistently appears at the wrong moment becomes noise. One that appears at exactly the

right moment begins to feel like expertise.

The Future of Product Experience

Traditional software became valuable because it allowed users to accomplish tasks. Intelligent software becomes valuable because it helps users accomplish those tasks better. That difference is subtle. It is also transformative. Consumers are gradually becoming accustomed to products that contribute ideas instead of merely executing instructions. As this expectation grows, products that simply wait for commands will increasingly feel passive.

The most successful products of the Intelligence Age will not compete solely on the quality of their responses. They will compete on the quality of their judgment. Their ability to notice what matters. To surface what is important. To recommend what users should consider next. In many ways, this marks the beginning of software behaving less like a tool and more like an experienced collaborator.

And great collaborators share another characteristic. They know that the value of intelligence is not measured by how much they say, but by how much unnecessary thinking they eliminate. That is the subject of the next principle.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Logprobs	Token-level probabilities. The cheapest available confidence signal, though a rough one.
Self-consistency	Sampling repeatedly and measuring agreement. More reliable than logprobs for judging whether a recommendation is stable.
Calibration	Aligning stated confidence with actual accuracy. Language models are systematically overconfident. This needs deliberate correction.
Abstention	Designing silence as a valid, well-handled output. The chapter's own line, that silence beats poor advice, requires this to be built, not hoped for.
LLM as judge	A model grading another model's output. Screen a recommendation before it reaches the user.
Guard models	Small classifiers checking safety or policy. A fast gate in front of anything proactive.
Reranking	Ordering candidates by predicted value. Generate several possible recommendations, then choose.
Tool use	The ability to act, not only advise. A recommendation acceptable in one step is worth more than one the user must execute.
Agent loop	Plan, act, observe. How an accepted recommendation becomes a completed action.
Preference data	Signals from what users accept and reject. The feedback making recommendation quality improve rather than freeze at launch.
Precision at k	Share of top suggestions that are good. The metric that matters when you only get to interrupt once.

Techniques

TECHNIQUE	WHAT IT MEANS
Generate several, show one	Produce candidates, rank them, surface the best.
Map confidence to behaviour	High recommends, medium offers alternatives, low asks, very low says nothing.
Always attach a reason	An unexplained recommendation cannot be evaluated by the person receiving it.
Cap interruption frequency	Recommendation fatigue erodes trust faster than being occasionally wrong.
Instrument accept and reject	Without this signal there is no path to improvement.

The Devil is in the TL;DR

One of the greatest misconceptions about artificial intelligence is that its value lies in generating more. More text. More ideas. More documents. More reports. More options. More information. In reality, the opposite is often true. The greatest value of intelligence lies not in producing more content, but in helping people consume less of it. Throughout the Information Age, information became abundant. Search engines, social media, messaging platforms, cloud storage, collaboration tools, and enterprise systems dramatically reduced the cost of creating and sharing information.

The result was extraordinary access to knowledge. It was also unprecedented information overload. Consumers are no longer limited by access to information. They are limited by attention. The scarce resource of the Intelligence Age is not information. It is cognition.

More Information Is Rarely Better

Imagine receiving a fifty-page report before an important meeting. Technically, all the information you need may be there. Practically, it is almost useless if you have fifteen minutes before the meeting begins. The problem is no longer information. The problem is extracting what matters. For decades, software largely ignored this responsibility. Documents became searchable. Emails became filterable. Dashboards became customizable.

But users remained responsible for reading, comparing, interpreting, and synthesizing the information themselves. Software improved access. Humans still performed understanding. Intelligent software changes that division of labour.

Compression Is Intelligence

One of the defining characteristics of human expertise is the ability to compress complexity. An experienced physician rarely explains every laboratory value individually. They summarize what matters. A senior engineer does not describe every implementation detail when reviewing a system. They identify the architectural trade-offs. An experienced consultant rarely presents every observation collected during research. They distill the three insights that actually change the decision.

Expertise is often measured by the quality of compression. Thoughtware should aspire to the same standard. Its role is not to expose every available fact. Its role is to surface the information that deserves human attention.

Every Product Is Becoming a Summarization Product

Most people associate summarization with documents. In reality, almost every product now faces the same challenge. An email application summarizes an inbox. A CRM summarizes customer history. A project management tool summarizes weeks of activity. A healthcare system summarizes patient records. A design platform summarizes stakeholder feedback. A financial application summarizes business performance. A recruitment platform summarizes candidate evaluations.

Even search itself is becoming an exercise in summarization rather than retrieval. Every product increasingly helps users answer the same question:

“What do I actually need to know?”

The products that answer this well reduce cognitive effort far more effectively than those that simply present more information.

The Right Summary Depends on the User

There is no universally correct summary. The information that matters to a CEO differs from what matters to an engineer. A physician reviewing patient history looks for different signals than a hospital administrator. A product manager preparing for a customer meeting needs different context than a sales representative. Good summarization is therefore deeply contextual. The objective is not to produce the shortest possible version.

It is to preserve the information that matters to the person making the decision. Intelligent products should increasingly adapt summaries based on role, expertise, goals, and the current situ-

ation. Summarization is personalization.

Summaries Should Lead to Decisions

One common mistake is treating summaries as compressed versions of original content. That is rarely their purpose. The purpose of a summary is to help someone decide what to do next. A meeting summary should make follow-up actions obvious. A financial summary should reveal where attention is required. A research summary should expose the conclusions that influence strategy.

A code review summary should identify the changes that deserve inspection. Every summary should answer an implicit question:

“Now that I know this, what should I do?”


Summarization without direction reduces reading. Summarization with direction improves decision-making. The difference is profound.

Progressive Detail

Reducing cognitive effort does not mean hiding information. It means revealing information progressively. Consumers should begin with the smallest amount of information necessary to understand the situation. Only then should they decide whether greater detail is required. An executive dashboard should begin with key insights before exposing detailed metrics. A legal document should begin with important clauses before presenting every page. (See Fig 25.)

A customer support system should surface the likely issue before displaying the full interaction history. Intelligent products should allow users to progressively expand their understanding rather

than forcing them to consume everything equally. The summary becomes the interface. The detail remains available when needed.



Good Intelligence Knows What to Omit

Perhaps the hardest part of summarization is deciding what not to include. Humans often assume intelligence is demonstrated by completeness. In reality, intelligence is frequently demonstrated by omission. The best teachers simplify without oversimplifying. The best writers remove unnecessary words. The best presenters eliminate unnecessary slides. The best advisors focus attention on what matters most. Intelligent software should follow the same principle.

Every additional sentence, chart, notification, or recommendation imposes another cognitive cost.

Products should therefore compete not by displaying more, but by requiring users to process less.

Designing for Cognitive Compression

Product teams have traditionally optimized for discoverability. The Intelligence Age introduces another responsibility. Compression. Where can the product reduce reading? Where can it reduce comparison? Where can it reduce searching? Where can it reduce decision fatigue? Where can it reduce context switching? These questions should become part of every product discussion. Summarization is not a feature. It is a philosophy of respecting human attention.

Intelligence Should Leave People Thinking Less, Not Reading More

Throughout the Information Age, software helped us find information. The Intelligence Age asks software to understand information before presenting it. Consumers increasingly expect products to tell them what matters rather than showing them everything. As this expectation grows, products that simply expose data without interpretation will begin to feel incomplete. The best intelligent products will quietly remove hundreds of small cognitive decisions every day.

Users may never notice how much thinking has been eliminated. They will simply notice that work feels easier. That is the hall-

mark of great intelligence. Not the amount of content it produces. The amount of unnecessary thinking it removes. Yet reducing thinking alone is not enough. Consumers also need confidence that the thinking being done on their behalf is sound.

Sometimes, the most reassuring thing an intelligent product can do is reveal how it arrived at its conclusion. That is the principle explored in the next chapter: Visible Thinking .

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Extractive summarization	Selecting sentences from the source. Cannot hallucinate, but reads poorly and misses synthesis.
Abstractive summarization	Generating new prose. Reads well, and can assert things the source never said.
Hierarchical summarization	Summarising in passes, then summarising summaries. How you handle documents larger than the context window, accepting loss at each level.
Context window	The attention budget. Determines whether a document can be summarised in one pass or must be chunked.
Groundedness	Whether claims are supported by the source. The property making a summary safe to act on.
Entailment checking	Automated verification that a claim follows from a passage. Machine-checkable faithfulness, cheap enough to run on every summary.
Citation and span attribution	Linking each claim to its origin. Converts a summary from assertion into something verifiable.
Structured extraction	Pulling named fields rather than prose. For contracts and invoices, extraction beats summary on accuracy and usefulness.
Document parsing	Layout-aware handling of tables, clauses, scans. A contract summary is only as good as the parse beneath it.
Vision-language models	Reading documents as images. Necessary for scans, complex layouts, anything where position carries meaning.
Reranking	Ordering by importance. What surfaces first is a ranking problem, not a writing problem.

Techniques

TECHNIQUE	WHAT IT MEANS
Salience	What counts as important, and to whom. A single canonical summary serves nobody especially well.
Eval frameworks	Tooling for faithfulness, relevance, coverage. Summarization is unusually well served by automated evaluation. Use it.
Time to first token	Delay before output starts. The summary must appear before the user gives up and scrolls.
Layer the disclosure	Headline, summary, key clauses, full source, letting the reader stop where they choose.
Cite everything	Every claim traceable to a span. No exceptions in high-stakes documents.
Extract fields for known types	When you know the document type, extract structure instead of writing prose.
Verify entailment automatically	Check faithfulness before display, not in review.
Vary salience by role	A lawyer and a buyer need different summaries of the same contract.

Visible Thinking

One of the biggest differences between traditional software and intelligent software is that traditional software rarely surprises us. When we click a button, we generally know what will happen. If we apply a filter, sort a table, or save a document, the software follows a deterministic sequence of operations. The process may be hidden, but it is predictable because the rules are explicit.

Thoughtware behaves differently. It interprets. It reasons. It weighs alternatives. It makes assumptions. It exercises judgment. That makes it dramatically more capable. It also makes it more opaque. When software begins participating in cognitive work, users naturally begin asking a new question:

“How did it arrive at that?”

The need for visible thinking emerges not because people are curious about artificial intelligence, but because they are trying to establish trust.

Humans Trust Reasoning More Than Answers

Imagine asking two consultants the same business question. The first replies,

“Expand into the European market.”

The second replies,

“I recommend expanding into Europe because your strongest customer segment already generates 28% of its traffic from that region, your competitors have limited presence there, and your existing logistics partners already support those countries. The primary risk is regulatory complexity, which can be mitigated by partnering locally during the first year.”

Both consultants may have reached the same conclusion. The second inspires significantly more confidence. Not because the answer is necessarily better. Because the reasoning is visible. Humans rarely trust expertise simply because someone claims to be right. They trust expertise because they can understand the thinking behind it.

Intelligence Is No Longer Just About Accuracy

During the Information Age, software was evaluated primarily on correctness. Did the calculation produce the right number? Did the file save successfully? Did the search return relevant results? Intelligent software introduces another dimension. Even when the answer is correct, users often want confidence that the reasoning itself is sound. This is especially true when decisions carry meaningful consequences.

A doctor reviewing a diagnosis. A lawyer reviewing a contract. A financial analyst reviewing an investment recommendation. A

hiring manager evaluating candidates. In these situations, the quality of reasoning matters as much as the quality of the conclusion. Products must therefore design not only for intelligence, but for confidence.

Thinking Is Part of the Experience


One of the most interesting behavioural shifts introduced by Large Language Models is that consumers have become surprisingly comfortable waiting. For decades, product teams fought relentlessly to reduce loading time. Every second of delay felt like friction. Then something unusual happened. People willingly watched language models think. The experience was objectively slower than many traditional interfaces. Yet it often felt faster.

Why? Because users understood that meaningful cognitive work was taking place. Consumers intuitively recognize that thinking requires time. They know that reviewing a legal contract, comparing twenty alternatives, or writing a thoughtful report cannot happen instantaneously. A brief period of visible reasoning often feels more trustworthy than an immediate answer. The waiting itself becomes evidence that the system is considering the problem rather than simply retrieving a predefined response.

Show the Thinking, Not Every Thought

Visible thinking is often misunderstood. It does not mean exposing every intermediate computation performed by the model. Users are rarely interested in the mechanics of inference. They care about the reasoning that influences the outcome. What assumptions were made? What evidence was considered? What trade-offs influenced the recommendation? What uncertainties remain? Why was one option preferred over another? (See Fig 26.)

The objective is explanation, not transparency for its own sake. Showing everything creates noise. Showing the right things creates understanding.



Visible Thinking Should Match the Decision

Not every interaction deserves the same level of explanation. If software corrects a spelling mistake, no explanation is necessary. If it recommends deleting a production database, an explanation becomes essential. The amount of visible thinking should increase with the importance, ambiguity, and consequence of the decision. Routine actions benefit from speed. Significant decisions benefit from understanding. This balance is one of the defining design challenges of intelligent products.

Too little explanation reduces trust. Too much explanation increases cognitive effort. The product should adapt its reasoning to the significance of the decision being made.

Thinking Creates Better Collaboration

One of the hidden benefits of visible thinking is that it enables better collaboration. When users understand the reasoning, they can improve it.

“Your assumption about our target audience isn’t correct.”

“That comparison ignores regulatory requirements.”

“You’ve optimized for cost, but reliability matters more here.”

These corrections become possible because the system has exposed its assumptions rather than hiding them. Human collaboration works in exactly the same way. Good colleagues rarely present conclusions without context. They explain how they arrived there. That explanation creates opportunities for refinement rather than blind acceptance. Intelligent software should encourage the same collaborative dynamic.

Confidence Matters as Much as Reasoning

Visible thinking should also communicate uncertainty. One of the defining characteristics of expertise is knowing the limits of one’s knowledge. An experienced consultant distinguishes between evidence and speculation. An experienced physician distinguishes between diagnosis and hypothesis. An experienced engineer distinguishes between known facts and informed assumptions. Intelligent software should behave similarly. Consumers are remarkably tolerant of uncertainty when it is communicated honestly.

They become distrustful when software expresses unwarranted confidence. Sometimes the most intelligent response is,

“I’m not certain.”

Or,

“There are two reasonable approaches, each with different tradeoffs.”

Reasoning without humility is often less trustworthy than uncertainty accompanied by good explanation.

Designing Visible Thinking

Designing visible thinking requires answering a different set of product questions. When should reasoning be shown? How much explanation is appropriate? Should users request explanations, or should they appear automatically? How should assumptions be presented? How should uncertainty be communicated? How should supporting evidence be referenced? These are experience decisions. The objective is not to prove that the model is intelligent.

The objective is to help users make better decisions with confidence. Visible thinking succeeds when it increases understanding without increasing cognitive burden.

Trust Is Built Before Accuracy Is Proven

Consumers rarely have the expertise to independently verify every recommendation an intelligent system makes. Instead, they rely on signals of credibility. Visible reasoning is one of the strongest such signals. It demonstrates that conclusions were reached through a process rather than produced as unexplained assertions. Over time, this changes the relationship between humans and software. Consumers no longer see the product as a machine that generates answers.

They begin to see it as a collaborator whose reasoning they can inspect, challenge, and improve. That relationship is essential to designing trustworthy intelligent systems. The Information Age rewarded products that responded quickly. The Intelligence Age rewards products that think well, and know when to let users see that thinking. Ultimately, however, even the most thoughtful reasoning cannot compensate for a product that delivers poor outcomes.

Consumers may admire elegant explanations, but they remain delighted by one thing above all else: Software that consistently helps them achieve better results. That is why, in the Intelligence Age, intelligence itself increasingly becomes the primary source of product delight.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Streaming	Token-by-token delivery. The simplest and most effective form of visible progress.
Time to first token	Silence before anything appears. The single largest driver of perceived slowness.
Tokens per second	Sustained output rate. Determines whether visible progress feels like progress or like struggle.
Reasoning models	Models spending extra compute before answering. They produce genuine intermediate work, but raw traces are rarely fit to display.
Chain of thought	Intermediate reasoning steps. Useful internally, but verbose, occasionally alarming, sometimes unfaithful as a user-facing artefact.
Explanation faithfulness	Whether a stated rationale is the actual cause. Be careful claiming the shown reasoning is why the system concluded what it did.
Tool-call traces	The record of actions taken. Showing what the system did is more honest and more useful than what it thought.
Agent loop	Discrete plan-act-observe steps. Gives you natural, nameable events to surface.
Tracing and spans	Structured instrumentation of execution. The same telemetry that debugs an agent can drive its progress display.
Checkpointing	Persisting intermediate state. Lets a long task survive interruption and resume visibly.
Durable execution	Long-running work outliving a session. Changes progress from a spinner into a status a user can leave and return to.

Techniques

TECHNIQUE	WHAT IT MEANS
Interruptibility	The ability to cancel mid-flight. Visible work implies stoppable work.
Semantic caching	Reusing answers for semantically similar requests. Instant responses need no progress display, but must not look suspicious.
Show doing, not thinking	"Reviewing 40 documents" beats a rendered monologue.
Name the step and count the units	Real status requires instrumentation, not animation.
Offer cancel wherever you show progress	Progress without control is frustration.
Summarise reasoning, never dump it	Condense the trace into something a person would say.
Reserve full reasoning for consequential output	Match depth of explanation to stakes of the decision.

Intelligence over Aesthetics

For much of the Information Age, the fastest way to improve a product was to improve its interface. Design teams refined typography, spacing, navigation, colour, motion, responsiveness, accessibility, and interaction patterns. These improvements transformed software from something functional into something enjoyable. Beautiful products stood out. Intuitive products won. Elegant interfaces became competitive advantages. Those advances fundamentally changed the software industry.

They remain just as important today. What has changed is not the importance of aesthetics. It is the relative importance of intelligence. Consumers increasingly forgive imperfect aesthetics when software behaves intelligently. They rarely forgive beautiful software that makes them think unnecessarily.

Delight Has Moved


Every era has a different source of delight. Early personal computers delighted people because they could perform calculations in seconds. The internet delighted people because information became instantly accessible. Smartphones delighted people because computing became portable. Today's consumers are delighted by something different. Software that understands them. Software that remembers. Software that saves them from repetitive thinking. Software that recommends better decisions. (See Fig 24.1)

Software that feels genuinely helpful. The emotional response no longer comes primarily from how software looks. It increasingly comes from how software behaves.

Where delight moved

FIG 24.1

As intelligence becomes table stakes, delight shifts from how things look to how intelligently they work.



Beauty Is Becoming Expected

Walk into a modern coffee shop almost anywhere in the world. You expect it to be clean. You appreciate thoughtful interior design. You notice comfortable seating. Yet none of these things are likely to become the primary reason you return. They are expectations. The quality of the coffee, the service, and the overall experience determine whether you become a regular customer.

Digital products are following the same trajectory. Consumers expect attractive interfaces. They expect smooth animations. They expect polished typography. They expect responsive interactions. Good aesthetics have become part of the baseline. Products increasingly differentiate themselves through intelligence layered on top of that foundation.

The Most Beautiful Interface Is Often the One You Barely Notice

One interesting consequence of intelligent software is that interfaces themselves often become smaller. When software understands natural language, remembers context, and recommends the next action, many traditional controls simply disappear. Products require fewer menus. Fewer settings. Fewer dialogs. Fewer workflows. The interface becomes quieter because intelligence absorbs complexity that previously had to be exposed visually. Paradoxically, intelligence often produces better aesthetics, not because designers focus more on beauty, but because they need fewer interface elements in the first place.

Good intelligence simplifies design.

Intelligence Creates Emotional Experiences

Traditional interface design has long focused on reducing friction. Intelligent products introduce another source of emotional value. Relief. The feeling that the software already understands. The satisfaction of seeing a thoughtful first draft instead of a blank page. The confidence of receiving a useful recommendation before making a difficult decision. The reassurance that previous conversations have been remembered.

These moments create emotional responses every bit as powerful as visual beauty. Sometimes more powerful. Consumers rarely describe these experiences by saying,

“The interface was elegant.”

They say,

“It just got me.”

That is a different kind of delight.

Great Design Is Becoming Less Visible

As products become more intelligent, much of the design work shifts away from what users can see. Designers spend less time deciding where a button belongs and more time deciding when software should make a recommendation. They spend less time organizing menus and more time designing memory. Less time exposing configuration and more time designing conversation. Less time arranging information and more time deciding what information should never need to be shown.

The centre of design gradually moves from pixels toward behaviour. Some of the most important design decisions become almost invisible. Users experience them without ever noticing they were designed.

Intelligence Does Not Excuse Poor Design

There is a temptation among AI products to assume that intelligence compensates for weak craftsmanship. It does not. An intelligent system hidden behind a confusing interface remains frustrating. A product with poor typography, inconsistent interaction patterns, inaccessible controls, or unclear information architecture still creates unnecessary cognitive effort. The Intelligence Age does not replace the craft of design. It raises the standard.

Products must now excel in both dimensions. Beautiful products that think well. Intelligent products that feel effortless. Neither alone is sufficient.

Designing for Behaviour Before Appearance

One useful exercise for product teams is to imagine evaluating their product without looking at the interface. Would users still find it valuable? Would its recommendations still improve decisions? Would its memory still strengthen the relationship? Would its reasoning still inspire confidence? Would it still reduce cognitive effort? If the answer is no, then the product's value depends largely on presentation rather than intelligence. That exercise is easier when intelligence is built as cognitive units, because a unit can be evaluated on its own. A product-level score tells you something is wrong. A unit-level score tells you where.

Conversely, imagine placing remarkable intelligence behind an average interface. Many users would continue returning because the product consistently improves their work. Behaviour increasingly creates loyalty. Appearance increasingly creates first impressions. Both matter. One is becoming more durable.

Intelligence Is Becoming the Brand

Historically, products became known for how they looked. People recognized operating systems, phones, and applications from their visual language. Increasingly, products are becoming known for how they think. One assistant gives concise answers. Another provides thoughtful reasoning. One remembers previous conversations. Another excels at research. Another plans projects particularly well. Their identities are increasingly defined by behaviour rather than appearance.

Consumers build relationships with products based on their intelligence. That relationship becomes part of the brand itself.

The New Source of Product Delight

The Information Age rewarded products that looked simpler than their competitors. The Intelligence Age rewards products that make users feel smarter than they would have without them. This is a profound shift. Consumers no longer judge products solely by how efficiently they execute commands. They judge them by how meaningfully they contribute to solving problems. A beautiful interface may attract attention.

Intelligent behaviour earns trust. A polished visual design may encourage someone to try a product. A thoughtful recommendation, a remembered preference, or a perfectly timed insight gives them a reason to return. The future of product design is therefore not a choice between aesthetics and intelligence. It is recognizing that aesthetics are becoming the foundation, while intelligence is becoming the primary source of delight.

The interface is no longer the whole experience. Increasingly, it is simply the window through which intelligence is experienced.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
Cognitive unit (CU)	A prompt with a defined input, a specified model, and an expected output. The smallest thing you can evaluate in isolation, which makes it the right unit to evaluate.
Eval sets	Curated cases with known good outcomes. The most important asset in an intelligent product. Build it before the feature.
LLM as judge	Model-based grading at scale. Makes evaluation affordable. Must be calibrated against human judgment on a sample.
Eval frameworks	Tooling for structured evaluation and metrics. Faithfulness, relevance, context precision are largely solved measurement problems now.
Offline vs online evaluation	Pre-release testing versus live measurement. They disagree more often than teams expect. You need both.
Regression testing	Re-running evals on every change. Model and prompt changes alter behaviour silently. Without this, quality drifts unnoticed.
Drift detection	Monitoring for behaviour change over time. Providers update models beneath you. Your product changes without a deploy.
Fine-tuning	Training on task-specific examples. Buys consistency and style, rarely knowledge. Frequently reached for too early.
LoRA and parameterefficient tuning	Adapting a model via a small trained overlay. Makes fine-tuning cheap enough to be a real option rather than a project.

Techniques

TECHNIQUE	WHAT IT MEANS
Preference tuning	Training on chosen-versus-rejected pairs. Turns accumulated user feedback into model behaviour.
Distillation	Training a small model on a large one's output. How quality survives a move to a cheaper, faster model.
Quantization	Reduced numerical precision. Cuts cost and latency, with quality effects you must measure rather than assume.
Model routing	Sending requests to different models by need. The main lever balancing cost against quality across a product.
Cost per interaction	Unit economics of inference. Intelligence carries marginal cost interface work never did. Now a design constraint.
Prompt versioning	Treating prompts as versioned artefacts. Prompts are production code. Untracked changes are untracked outages.
Observability	Tracing, logging, inspection of live behaviour. You cannot improve what you cannot see in production.
Build the eval set first	Before the feature. It defines what good means.
Calibrate the judge	Check model grading against human grading on a sample, regularly.
Re-run on every model change	Including provider updates you did not initiate.
Choose the smallest model that passes	Capability you do not need is cost and latency you do pay.
Track cost per successful outcome	Not cost per call. Failures are expensive twice.

A New Interaction Model

The first generation of computers was programmed through code. The next generation was operated through graphical interfaces. The web introduced hyperlinks and navigation. Mobile devices popularized touch. Voice assistants brought speech into mainstream computing. Each era introduced a new way for humans to communicate with machines. The Intelligence Age is different. It is not introducing another interface. It is introducing another layer of interaction.

Instead of learning the language of software, users increasingly expect software to understand the language of humans. This final theme explores how interaction itself is changing as intelligence becomes the primary medium between people and software.

Language over Layout

For decades, software interfaces were designed around layout. Every capability needed a place. Buttons occupied toolbars. Menus organized commands. Tabs grouped functionality. Forms collected inputs. Dashboards displayed information. Designers spent countless hours deciding where each capability should live so that users could eventually learn where to find it. The interface became a map. Users became navigators. The better someone knew the map, the more productive they became.

This model has shaped nearly every digital product we use today. Intelligent software introduces a radically different possibility. Instead of navigating the map, users increasingly communicate what they want through language. The interface becomes less about location and more about conversation.

Language Is Humanity's Native Interface

Humans begin speaking years before they learn to read. Long before children understand menus, forms, or icons, they understand conversation. Language is the oldest and most natural interface humanity has ever developed. For thousands of years, it has been our primary mechanism for sharing ideas, coordinating work, teaching, negotiating, questioning, and solving problems together. Ironically, computers required us to abandon that interface.

We learned clicks instead of conversation. Menus instead of meaning. Commands instead of intent. The graphical user interface was a remarkable improvement over programming. But it still required humans to adapt themselves to machines.



*For the first time in computing history,
machines are beginning to adapt themselves
to human language instead.*

Layout Exists Because Software Could Not Understand

Why do applications have hundreds of buttons? Because software historically lacked understanding. Every capability had to be explicitly exposed. If users wanted to perform an action, they first had to discover where that capability existed within the interface. The more powerful the application became, the more complicated the layout became. This relationship was unavoidable. Complex

capability demanded complex interfaces.

Language changes that trade-off. If software understands the request,

“Create a quarterly sales dashboard comparing regional growth over the last two years,”

the user no longer needs to locate dozens of individual visualization tools, formatting options, filtering controls, and layout settings. The complexity still exists. It simply moves behind intelligence.

Language Is Universal

One of the most remarkable properties of language is that it works across almost every medium. A user can describe a document. An image. A video. A spreadsheet. A presentation. A website. A piece of code. A financial model. A legal contract. The interaction remains remarkably consistent. Only the underlying artifact changes. Traditional interfaces often required completely different workflows for each medium.

Language creates a common interaction layer above them all. This universality is one of the reasons Large Language Models feel so transformative. They are not merely solving individual problems. They are creating a common interface for interacting with almost any digital artifact.

Language Is Not Just Typing

When people hear “natural language interface,” they often imagine a chat box. That is far too narrow a definition. Language increasingly combines with every other interaction model. Users highlight text and describe changes. They point at interface elements and explain what should happen. They upload screenshots and ask questions. They sketch rough ideas and describe the desired outcome.

They speak while manipulating physical objects. The interaction becomes multimodal. Language acts as the connective tissue between different forms of input. The future is not conversational software replacing graphical interfaces. It is language enriching every interface.

Layout Does Not Disappear

The rise of language does not make interfaces obsolete. There will always be actions that are faster through direct manipulation. Adjusting a volume slider. Dragging an object. Scrubbing through a video. Playing a song. Accepting a notification. Viewing information. Graphical interfaces remain exceptionally effective for deterministic interactions. The opportunity lies elsewhere. Whenever users need to express ambiguity, describe intent, explore ideas, or perform complex cognitive tasks, language becomes dramatically more efficient than navigating interfaces. (See Fig 27.)

The future belongs to products that seamlessly combine both interaction models. Direct manipulation where precision matters. Language where understanding matters.

What language takes, what layout keeps

FIG 27

Layout does not disappear

 LANGUAGE TAKES When understanding is required	 LAYOUT KEEPS When doing is enough
 Expressing intent I want, I need, I prefer...	 Volume slider Set level, adjust precisely.
 Ambiguity When the task or information is not fully specified.	 Dragging an object Move, resize, reposition.
 Multi-step requests Do this, then that, and also...	 Scrubbing video Seek, scan, fine control.
	 Playing a song Play, pause, next, previous.
 Anything needing understanding Judgment, reasoning, context, trade-offs, priorities.	 Accepting a notification Approve, dismiss, mark as read.
	 Viewing information Scan, compare, read.

Language Changes Discovery

Traditional software assumes that users discover capabilities by exploring interfaces. Language introduces a different model. Users discover capability simply by asking. Many consumers have already experienced this shift. Instead of wondering whether a feature exists, they ask.

“Can you generate a summary?”

“Can you convert this into a presentation?”

“Can you compare these contracts?”

“Can you prepare interview questions?”

The product responds either by completing the task or explaining its limitations. This dramatically reduces the learning curve. Con-

sumers no longer need to memorize what software can do. They simply express what they hope it can do.

Designing Language-First Products

Designing for language requires a different philosophy. Interfaces should help users express intent rather than merely expose functionality. Products should gather supporting context automatically before asking unnecessary questions. Suggestions should help users discover possibilities without requiring them to study documentation. Most importantly, language should complement, not compete with, traditional interaction patterns. Every task should use the interaction model that minimizes cognitive effort.

Sometimes that will be a button. Sometimes a gesture. Sometimes a sentence. The role of the product is to make that distinction invisible.

The Universal Interface

Throughout computing history, every new interface has reduced the distance between human thought and machine execution. Programming required people to think like computers. Graphical interfaces allowed computers to represent familiar physical metaphors. Touch interfaces removed intermediary devices such as keyboards and mice. Language removes something even more fundamental. The need to translate thought into interface operations. Consumers increasingly expect software to understand objectives expressed in ordinary language.

As this expectation spreads, layout becomes less about exposing capability and more about supporting understanding. The interface no longer serves primarily as a control panel. It becomes a

collaborative workspace where language, visual interaction, documents, images, code, and other media coexist naturally. The future of interaction is therefore unlikely to be defined by the disappearance of layout. It will be defined by the rise of language as the universal layer that connects every layout, every modality, and every form of digital work.

Only one question then remains. Once software understands us, remembers us, reasons with us, and communicates naturally, how should humans and intelligent systems actually make decisions together? That is the final principle of this book.

What Makes This Work

Concepts

CONCEPT	WHAT IT MEANS
ASR	Speech to text. The entry point for voice. Accuracy matters less than latency and recovery.
TTS	Text to speech. Output quality shapes trust more than most teams anticipate.
Speech-to-speech models	Audio in, audio out, without a text hop. Removes the latency making voice interfaces feel like walkietalkies.
Voice activity detection and endpointing	Knowing when a person started and stopped speaking. Bad endpointing is the most common cause of a voice interface feeling broken.
Barge-in	Letting the user interrupt playback. Non-negotiable for anything conversational.
Round-trip latency budget	Total time from speech to response. Voice is far less forgiving than text. The tolerable ceiling is a fraction of a second.
Multimodal models	Text, image, and audio in one system. Lets language refer to what the user is looking at.
Vision-language models	Models that read images. "Fix this bit" requires the system to see the bit.
Function calling	Utterance to typed action. Where language interfaces most often break, not in understanding, but in resolving to something executable.
Constrained decoding	Restricting output to a valid grammar. For high-stakes commands, a bounded vocabulary remains more reliable than open language.

CONCEPT	WHAT IT MEANS
Multilingual models	Coverage across languages, including mixed input. Language interfaces expose coverage gaps immediately, including code-switching mid-sentence.

Techniques

TECHNIQUE	WHAT IT MEANS
Semantic caching	Matching new requests to previously answered similar ones. Common utterances can return instantly.
Word error rate	Standard ASR accuracy metric. Useful, but secondary. Task success per utterance is the metric that matters.
Resolve every utterance to a typed action	Understanding without execution is a demo.
Constrain the vocabulary for highstakes actions	Know where to trade expressiveness for reliability.
Design the repair path first	Misheard input is the normal case, not the edge case.
Teach affordances by example	Users cannot see what they may say. Show, suggest, seed.
Keep a non-language alternative	Language should be the primary interface, not the only one.

Co-pilot the Decision

The ultimate promise of the Intelligence Age is not automation. It is collaboration. For decades, software has existed at one of two extremes. At one end, software behaved like a passive tool. It waited for instructions, executed them precisely, and stopped. Every decision belonged to the user. At the other end, automation attempted to remove the user entirely. Rules were defined in advance, workflows executed automatically, and humans intervened only when something went wrong.

Thoughtware introduces a third model. Neither passive tools nor fully autonomous systems. Collaborators.

The future of product design lies not in replacing human judgment, but in augmenting it.

The Wrong Question

Much of the conversation around AI asks a single question:

“Will AI replace humans?”

For product builders, that is rarely the most useful question. A better one is:

“Which parts of this decision should belong to the human, and which should belong to the intelligence?”

Almost every meaningful product decision can be divided into smaller cognitive tasks. Gathering information. Identifying patterns. Generating alternatives. Estimating consequences. Comparing trade-offs. Making a final judgment. Not all of these tasks require the same kind of intelligence. Some are repetitive. Some require expertise. Some require values. Some require accountability. The objective is not to automate the entire decision.

It is to distribute the decision intelligently.

Humans and Machines Have Different Strengths


Intelligent systems excel at processing large amounts of information, identifying patterns, maintaining consistency, recalling previous knowledge, and exploring many alternatives quickly. Humans excel at understanding nuance, interpreting social context, exercising moral judgment, balancing competing values, and accepting responsibility for consequences. These strengths are complementary rather than competitive. A hiring system may compare hundreds of candidates objectively. A hiring manager understands team dynamics that never appear in a résumé. (See Fig 28.)

A financial assistant may identify unusual spending patterns. A business leader understands strategic priorities that exist outside the data. A medical system may recognize statistical probabilities. A physician understands the emotional and personal context surrounding treatment. The strongest decisions emerge when both perspectives contribute.

Complementary strengths, meeting at the decision

FIG 28

Humans and machines have different strengths



The Decision Is a Conversation

Traditional software treated decisions as transactions. Users collected information. Software executed the final action. Thoughtware increasingly transforms decisions into conversations. The system proposes. The user questions. The system explains. The user introduces additional constraints. The system revises. Together they gradually converge on a better outcome. This re-

sembles how experienced colleagues work together. Few important decisions emerge perfectly formed.

They improve through discussion. Intelligent products should encourage the same process. The objective is not immediate agreement. It is progressive refinement.

Intelligence Should Challenge, Not Just Agree

One of the easiest ways to make intelligent software feel helpful is to make it agreeable. Unfortunately, agreement is not the same as value. The best colleagues occasionally disagree. They notice risks. Question assumptions. Offer alternative perspectives. Challenge weak reasoning. Prevent costly mistakes. Consumers increasingly expect intelligent software to do the same. If a product blindly accepts every proposal, it contributes little beyond execution.

If it respectfully challenges questionable decisions, it becomes a genuine collaborator. This requires balance. The goal is not to argue. It is to improve thinking. Great collaborators know when to support and when to challenge. Intelligent products should learn the same habit.

Preserve Human Agency

As intelligent systems become more capable, one design principle becomes increasingly important. The user should always understand who is making the decision. Recommendations are valuable because they preserve agency. Automation is valuable because it removes repetitive effort. Confusion arises when products blur the boundary between recommendation and decision. Users should know whether the system is suggesting an action, request-

ing approval, or acting independently.

The distinction should never be ambiguous. Trust depends upon clarity. Consumers are generally comfortable delegating routine decisions. They become significantly more cautious when consequences become meaningful. Good product design respects that instinct.

Confidence Should Determine Autonomy

Not every task deserves the same level of automation. A system may confidently rename files without asking. It should probably not approve a multimillion-dollar acquisition without review. The appropriate level of autonomy depends on three factors. The confidence of the system. The consequence of failure. The reversibility of the decision. When confidence is high, consequences are small, and mistakes are easily reversed, software should increasingly act autonomously. (See Fig 29.)


As any of these factors change, human involvement should increase proportionally. Autonomy is not binary. It exists on a spectrum. One of the defining responsibilities of Intelligence Age product design is deciding where each interaction belongs on that spectrum.

The autonomy spectrum

FIG 29

Autonomy is not binary. It's a spectrum.

Placing interactions on it is a defining responsibility of the discipline.



Collaboration Improves Both Participants

One interesting characteristic of human collaboration is that both participants improve. People teach each other. Correct each other. Adapt to each other's preferences. Develop shared language. Build trust. Intelligent products should aspire to the same relationship. The software should gradually learn the user's preferences, priorities, and working style. Equally, users should develop confidence in the strengths and limitations of the software.

Over time, the relationship becomes more efficient because both sides understand each other better. The experience shifts from using a product to working with a partner.

Designing Collaborative Systems

Designing collaborative intelligence requires product teams to think beyond individual interactions. How should disagreements be handled? When should software insist on a recommendation? When should it simply offer an alternative? How should conflicting objectives be surfaced? How should responsibility be communicated? How should decisions be revisited as new information becomes available? These questions extend far beyond interface design.

They concern the nature of collaboration itself. The most successful intelligent products will not simply provide answers. They will help users think better. Sometimes that means generating ideas. Sometimes summarizing complexity. Sometimes making recommendations. Sometimes challenging assumptions. Sometimes remaining silent. Knowing which role to play is itself an act of intelligence.

From Operator to Director

The central argument of this book has been remarkably simple. Consumers are transitioning from operating software to directing intelligence. This final principle completes that transition. A director does not perform every task personally. Nor does a director disappear from the process. A director establishes objectives, provides guidance, reviews progress, makes judgments, and accepts responsibility for the outcome. That is increasingly becoming the role of the user.

Intelligent software gathers information, generates alternatives, performs routine cognitive work, and executes appropriate actions. The human provides purpose. The intelligence provides

capability. Together, they produce outcomes that neither could achieve as effectively alone.

Designing for the Intelligence Age

Every age eventually develops a design philosophy that reflects the capabilities of its time. The Information Age taught us to design interfaces that were usable, beautiful, and efficient. The Intelligence Age asks something more. It asks us to design relationships. Relationships built on understanding instead of commands. Memory instead of repetition. Context instead of explanation. Recommendations instead of passive responses.

Reasoning instead of mystery. Conversation instead of configuration. Collaboration instead of operation. The sixteen chapters in this book have explored these shifts from different perspectives, but they all emerge from the same underlying behavioural change. Consumers increasingly expect software not merely to execute work, but to participate in it. The products that define the next generation will not simply contain artificial intelligence.

They will be designed around intelligent collaboration. The future of software is not human or machine. It is human with machine. That is what it means to design for the Intelligence Age.

What Makes This Work


Concepts

CONCEPT	WHAT IT MEANS
Agent loop	Plan, act, observe, revise. The execution model behind anything that acts rather than answers.
Autonomy levels	A graded scale from suggesting to acting silently. The spectrum the chapter argues for, expressed as system configuration.
Human in the loop	Designed points requiring human authorisation. Where the approval gates sit, and what they gate.
Tool use	The model's ability to take real action. Makes autonomy meaningful and dangerous in equal measure.
MCP	Standard protocol for tool and data access. Broadens what an agent can reach, which broadens what it can get wrong.
Least privilege and scoped credentials	Granting only the access a task requires. The most effective single control on agent blast radius.
Sandboxing	Execution against an isolated replica. Lets consequential plans be rehearsed before they are real.
Dry run	Producing the plan and its effects without committing. The preview making delegation acceptable.
Checkpointing and rollback	Saved states an action can be reversed to. Reversibility is engineered, not assumed. The cheapest way to buy autonomy.
Guardrails	Policy enforced outside the model. Rules the system cannot reason its way past.

CONCEPT	WHAT IT MEANS
Prompt injection	Instructions hidden in content the model reads. An agent that acts on what it reads can be hijacked by what it reads.
Skill	A known procedure that performs a capability. Package what is settled so the agent spends its uncertainty where uncertainty remains.
Acting skill	A procedure that changes the world without owning a goal. Acting is not agency, so do not promote it to one.
Path stabilisation	Naming a route an agent keeps rediscovering. Turns repeated deliberation into a capability.

Techniques

TECHNIQUE	WHAT IT MEANS
Indirect prompt injection	Injection arriving via retrieved documents, web pages, or email. The realistic attack path once an agent has both retrieval and tools.
Specification gaming	Satisfying the stated goal while missing the intent. Why trajectory matters as much as outcome.
Confidence estimation	Assessing how sure the system is. One of the three gates determining where an action sits on the spectrum.
Trajectory evaluation	Grading the path, including when it stopped. Tests whether the system knows when to hand back.
Audit logging	Immutable record of action and authority. What makes autonomy defensible in any regulated setting.
Gate by consequence and reversibility	Not by capability. What the system can do is the wrong question.
Make undo cheap before making autonomy wide	Every increment of reversibility buys an increment of trust.
Rehearse in a sandbox	Assess judgment before granting reach.
Log authority, not just action	Who authorised it, on what basis, under which policy.
Test for knowing when to stop	The hardest capability to build, the easiest to omit from evaluation.



PART III

Applying the Philosophy

Every design philosophy eventually faces the same question.

“This makes sense in theory. How do we actually build it?”

The previous chapters argued that the Intelligence Age is fundamentally changing consumer behaviour. They introduced sixteen principles that emerge naturally from that shift and explored how intelligence reshapes the relationship between people and software. Understanding those principles is only the beginning. Applying them is considerably more difficult. Unlike previous technological transitions, the Intelligence Age is not introducing a single new interface or a single new technology. It is changing every layer of a product simultaneously: its strategy, architecture, interaction model, information flow, business processes, and even the way product teams themselves operate.

Designing an intelligent product is therefore not a matter of adding an AI assistant or integrating a language model. It requires rethinking the product from first principles. Many organisations are currently approaching AI as a feature. A chatbot is added to the support page. A content generator is added to the editor. A summarization button is added to reports.

These features often provide genuine value. But they rarely change the product. The products that will define the Intelligence Age will not simply have AI features. They will be built around intelligence from the very beginning. That requires a different design process. This final part of the book explores how product teams can apply the philosophy presented in the earlier chapters, not as isolated features, but as a way of thinking about products.

Designing Products for the Intelligence Age

Every generation of product builders inherits assumptions from the generation before it. For decades, product teams have followed a remarkably consistent process. Identify user problems. Map user journeys. Design workflows. Create interfaces. Build features. Measure engagement. Improve usability. This process produced extraordinary products. It also reflects the assumptions of the Information Age. The user performs the thinking. The software performs the execution.

The interface becomes the centre of the product. The Intelligence Age changes those assumptions. The question is no longer simply,

“How do users complete this workflow?”

It becomes,

“How should intelligence participate in this workflow?”

That single question changes almost every design decision.

Start with Cognitive Work

Traditional product design often begins with tasks. What steps does the user perform? Which screens are required? How does information move through the application? Thoughtware design should begin one level higher. Where does the user spend the most cognitive effort? Where do they repeatedly make similar decisions? Where do they search for information? Where do they compare alternatives?

Where do they translate intent into procedure? Where do they repeatedly explain the same context? These moments represent opportunities for intelligence. The objective is not to automate workflows. It is to reduce unnecessary thinking. A useful exercise is to examine every step of a user journey and ask a simple question:

“Why does the human have to think here?”

If the answer is simply because software has always required it, that step deserves to be challenged.

Design Intelligence Before Interfaces

Traditional design naturally begins with screens. Intelligent products benefit from a different sequence. Before drawing a single interface, product teams should define how intelligence behaves. What should it know? What should it remember? What should it recommend? When should it ask questions? When should it remain silent? When should it act autonomously? Only after these behaviours have been designed should interfaces be created to support them.

This is a subtle but profound shift. The interface becomes an expression of intelligence rather than the primary product itself.

Treat Context as Infrastructure

Many intelligent products fail because they focus on models while ignoring context. A powerful model without context behaves like a knowledgeable stranger. It may produce intelligent responses, but it lacks the situational understanding necessary for meaningful collaboration. Every product should therefore identify the contextual signals available during each interaction. Current activity. Previous conversations. Selected content. Project history. Organisational knowledge.

Role. Objectives. Collaborators. Time. Location. The richer the context, the less effort users spend explaining themselves. Context is no longer supporting information. It is infrastructure.

Design Memory Deliberately

Memory should never emerge accidentally. It should be intentionally designed. Product teams should define what deserves temporary memory, what deserves long-term memory, and what should never be remembered at all. Users should understand what the product remembers and remain in control of that relationship. Memory should strengthen trust rather than weaken it. The best memory systems quietly eliminate repetition without becoming intrusive.

They make every future interaction feel more personal than the last.

Think in Recommendations, Not Features

Traditional roadmaps often describe functionality. Add dashboard. Add export. Add notifications. Add templates. Intelligent products benefit from another perspective. What decisions should this product improve? What recommendations should it make? When should those recommendations appear? How will users respond to them? This shifts attention away from capabilities and toward outcomes. The product no longer succeeds because it contains many features.

It succeeds because it consistently improves the user's decisions.

Design Behaviour, Not Responses

Many early AI products focused almost exclusively on response quality. Can the model answer correctly? Can it write naturally? Can it generate convincing text? These questions remain important. They are no longer sufficient. Consumers experience behaviour rather than responses. Does the product remember appropriately? Does it know when to ask for clarification? Does it challenge poor assumptions? Does it explain important decisions?

Does it recover gracefully from mistakes? Does it understand when not to interrupt? These behaviours define the personality of the product. Over time, they become far more memorable than individual responses.

Build Systems, Not Features

Perhaps the most important shift for product teams is recognising that intelligence is rarely a single feature. It is a system. Language models. Memory. Retrieval. Planning. Tools. Skills. Agents. Evaluation. Guardrails. Feedback loops. Each contributes to the overall experience. Removing any one of them changes the behaviour of the product. Great intelligent products are therefore designed as ecosystems rather than isolated capabilities. In practice that ecosystem has a grain, though not a strict one. Cognitive units decide, skills perform, agents pursue, and they compose freely rather than nesting in a fixed order.

Product leaders must think architecturally. Not because they are writing code, but because architecture increasingly shapes experience.

Measure Cognitive Outcomes

Traditional analytics answer questions about activity. How many users? How many sessions? How many clicks? How long did they stay? Intelligent products require another set of metrics. How much thinking was eliminated? How often were recommendations accepted? How frequently did users edit generated drafts? How much context was automatically gathered? How often did memory prevent repetition? How much time was saved?

Did users make better decisions? These metrics align much more closely with the actual purpose of intelligent products. The objective is not engagement for its own sake. It is cognitive leverage.

Build for Trust

Every intelligent capability should strengthen trust. Can users understand why a recommendation appeared? Can they inspect memory? Can they override decisions? Can they recover from mistakes? Can they confidently delegate routine work while retaining meaningful control? Trust should never be treated as an afterthought. It is the foundation upon which every intelligent interaction depends.

Without trust, even the most capable intelligence remains unused.

The New Product Mindset

Designing for the Intelligence Age ultimately requires a different way of thinking. Not every problem requires AI. Not every interaction benefits from conversation. Not every decision should be automated. The objective is not to maximise intelligence. It is to apply intelligence where it meaningfully improves human capability. Great products have always respected human attention. The next generation of products must also respect human cognition. (See Fig 30.)

That is the true opportunity of the Intelligence Age. Not replacing human intelligence. Amplifying it. The final chapter looks beyond today's products and considers where this philosophy may lead over the coming decades, and what responsibilities product builders carry as intelligence becomes an increasingly fundamental part of everyday life.

FIG 30

The Intelligence Age design canvas

A PRACTICAL FRAMEWORK FOR BUILDING INTELLIGENT EXPERIENCES

Use this canvas to plan, design, and evaluate intelligent systems that augment human cognition and earn trust over time.

Project / Initiative _____

Team _____

Date _____

Version _____



1 START WITH COGNITIVE WORK

What cognitive work are we helping people do?

Primary cognitive jobs:

Who does it today?

Where does it break down?

Why does it matter?



2 DESIGN INTELLIGENCE BEFORE INTERFACES

What should the system think and decide?

Decisions the system should make:

Judgments it should support:

Boundaries (what it should not do):

Inputs needed to think well:



3 TREAT CONTEXT AS INFRASTRUCTURE

What context does the system need, and how will we get it?

Context types (user, task, environment, domain):

Sources of context:

How context is captured:

How it stays current:

How it's used:



4 DESIGN MEMORY DELIBERATELY

What should the system remember, and for how long?

What to remember (facts, preferences, history):

Granularity & structure:

Retention rules (TTL / decay):

User control (view, edit, delete):

Privacy & security:



5 THINK IN RECOMMENDATIONS

What should the system suggest, and why?

Recommendations to make:

When to be proactive vs. passive:

Confidence signals to show:

Alternatives to present:

When to ask for clarification:



6 DESIGN BEHAVIOUR

How should the system behave in different situations?

Autonomy level (see spectrum): _____

Triggers to act:

Escalation / handoff rules:

Failure handling & safe defaults:

Tone, style & personality:



7 BUILD SYSTEMS, NOT FEATURES

What components work together to deliver intelligence?

Core components (models, tools, data, rules):

How they connect:

Feedback loops:

Human in the loop (where & how):

Scalability & reliability considerations:



8 MEASURE COGNITIVE OUTCOMES

How will we know we're improving how people think and decide?

Outcomes that matter:

Metrics (quantitative):

Signals (qualitative):

Baselines & targets:

How we'll learn and iterate:



9 BUILD FOR TRUST

How will we earn, communicate, and protect trust?

Transparency (what to explain, when):

Control (user choices & overrides):

Reliability (accuracy, consistency, robustness):

Fairness & bias considerations:

Accountability (who is responsible):



10 SHIP, LEARN, EVOLVE

How will we continuously improve intelligence and experience?

Release approach:


What we'll monitor:

How we'll incorporate feedback:

Hypotheses to test next:

Long-term vision:

AUTONOMY SPECTRUM (WHERE THIS INTERACTION FITS)



SYSTEM CONFIDENCE

Expected confidence range:

Low _____ High _____

How confidence is estimated / shown:

CONSEQUENCE OF FAILURE

Potential impact:

☐ Low

☐ Moderate

☐ High

☐ Severe

Notes: _____

REVERSIBILITY OF ACTION

How easily can the action be undone?

☐ Very easy

☐ Easy

☐ Difficult

☐ Very difficult

Notes: _____



THE DISCIPLINE'S RESPONSIBILITY

Continuously assess these factors, place interactions intentionally on the spectrum, and design for outcomes that enhance human capability and dignity.

NEXT STEP

Beyond the Interface

Every technological age eventually reaches a point where its defining technology becomes invisible. We rarely think about electricity while switching on a light. We rarely think about the internet while sending a message. We rarely think about databases while transferring money. The technology does not disappear. It simply fades into the background. Artificial intelligence will follow the same path.

Within a few years, consumers will stop asking whether a product uses AI. The question itself will begin to feel strange. Just as nobody advertises that their application “uses databases,” intelligent behaviour will simply become part of the expectation. The winners of the Intelligence Age will not be the products with the most visible AI. They will be the products whose intelligence feels natural enough to disappear.

Intelligence Will Become Infrastructure

Today, many intelligent products are experienced as destinations. People open a chatbot. Ask questions. Generate content. Close the application. This is unlikely to remain the dominant interaction model. Intelligence will increasingly become infrastructure embedded inside every product rather than a destination in itself. Every document editor will think. Every CRM will reason. Every healthcare system will recommend. Every design tool will collaborate.

Every development environment will plan. Every enterprise application will understand language. Consumers will stop distinguishing between software and AI-powered software. There will simply be software. And it will be intelligent.

Interfaces Will Become Quieter

As intelligence moves deeper into products, interfaces themselves will continue becoming simpler. Much of today's interface complexity exists because software requires explicit instructions. Buttons expose functionality. Menus expose capability. Settings expose behaviour. As software becomes increasingly capable of understanding intent, remembering context, and making recommendations, many of these controls will quietly disappear. Interfaces will not vanish. They will become calmer. (See Fig 31.)


Less occupied with exposing capability. More focused on supporting understanding.

The greatest compliment future interfaces may receive is that users hardly notice them.

Interface get quieter

As systems become more capable, the interface shrinks.
Less to operate. More to understand. Easier to trust.

FIG 31



Products Will Be Judged Like People

One fascinating consequence of intelligent behaviour is that consumers increasingly describe software using human language.

"It understands me."

"It forgot what I told it."

"It keeps interrupting."

"It gave me a thoughtful suggestion."

"It explains things well."

These are not descriptions of interfaces. They are descriptions of relationships. As products become more intelligent, they will increasingly be evaluated using the same qualities we value in human collaborators. Competence. Reliability. Judgment. Memory. Honesty. Humility. Empathy. Consistency. Trustworthiness. This does not mean products become human. It means the nature of interaction becomes increasingly social rather than mechanical.

That is a profound change for product builders.

Every Product Becomes a Relationship

For decades, software experiences were largely transactional. Open the application. Complete the task. Close the application. The next interaction began almost independently of the last. Thoughtware increasingly behaves differently. Every interaction influences the next. Memory accumulates. Understanding deepens. Recommendations improve. Preferences evolve. Trust develops. Products become long-term collaborators rather than temporary tools. The quality of that relationship will become one of the strongest competitive advantages any company can build.

Consumers rarely remain loyal to interfaces. They become loyal to relationships.

Product Design Becomes Behaviour Design

Throughout this book, one theme has repeatedly emerged. The role of the designer is expanding. Traditionally, designers shaped appearance. Then they shaped interaction. Increasingly, they shape behaviour. How should the product think? What should it remember? When should it speak? When should it remain silent? When should it recommend? When should it challenge? How

should it recover from mistakes?

How should it earn trust? These questions are no longer secondary considerations. They are rapidly becoming the centre of product design. Future design systems may include not only colours, typography, spacing, and components, but also behavioural patterns, memory strategies, recommendation principles, confidence guidelines, and collaboration models. The design system itself becomes intelligent.

The Responsibility of Product Builders

Every technological revolution expands human capability. It also expands human responsibility. The products we build increasingly influence how people learn, work, make decisions, solve problems, and collaborate with one another. Intelligence amplifies those effects. A thoughtful recommendation can save hours of work. A poor recommendation can quietly influence thousands of decisions. A well-designed memory system can strengthen trust.

A poorly designed one can destroy it. Product builders therefore carry responsibilities that extend beyond usability. We are increasingly designing systems that participate in human thinking. That responsibility demands humility. The objective should never be replacing human judgment. It should be strengthening it.

The Philosophy in One Sentence

Every chapter in this book has explored a different aspect of the same transformation. Low-Cognitive First Draft reduces the effort of beginning. Point & Fix removes unnecessary editing complexity. Intent over Instructions shifts attention from procedure to outcomes. Conversation over Configuration allows products to

adapt through dialogue. Context and Memory enable understanding across moments and across time. Recommendations improve judgment.

Summarization respects attention. Visible Thinking builds trust. Intelligence over Aesthetics recognizes where delight is moving. Language over Layout changes how humans communicate with software. Co-pilot the Decision redefines collaboration itself. Although each principle addresses a different aspect of product design, they all emerge from one underlying behavioural change. Consumers are transitioning from operating software to directing intelligence. Everything else follows from that observation.

The Beginning, Not the End

This book does not attempt to predict every interaction pattern that will emerge over the coming decades. Many of today's assumptions will undoubtedly change. New models will appear. New interfaces will emerge. New technologies will replace existing ones. Some principles described here may evolve. Others may become so obvious that future generations wonder why software was ever designed differently.

That is the nature of every technological transition. The goal of this book has never been to document today's tools. It has been to describe the behavioural shifts that those tools have set in motion. Technologies evolve. Human expectations evolve more slowly. Design philosophies endure longer still. If the ideas in these pages help product managers ask better questions, help designers think beyond interfaces, help engineers appreciate the behavioural consequences of architecture, or help founders build products that genuinely respect human cognition, then this book

will have achieved its purpose.

Designing the Next Age

The Information Age challenged us to make software usable. The Intelligence Age challenges us to make software thoughtful. That is a much harder problem. It is also a much more meaningful one. For the first time in history, software is becoming capable of participating in the cognitive work that once belonged almost exclusively to humans. Whether that future becomes empowering or overwhelming will depend less on the intelligence of our models than on the wisdom of our products.

That is what thoughtware means. Not software that answers, but software that thinks well enough to be trusted with the thinking.

Technology gives us capability. Design determines how that capability shapes human life. That has always been true. It is simply becoming more important than ever. The Intelligence Age has already begun. The question is no longer whether we will design thoughtware. The question is what kind of intelligence people will experience every day through the products we choose to build.

That future is not being written by technology alone. It is being designed.

Glossary

A quick reference to the ecosystem of thoughtware: the concepts that explain how these systems work, and the techniques that decide whether they work well. Each entry gives a one-line meaning and the chapter and page where the term is used, shown as chapter · page.

A

Abstention	Designing silence as a valid, well-handled output. 11 · 129
Abstractive summarization	Generating new prose. Reads well, and can assert things the source never said. 12 · 138
Acting skill	A procedure that changes the world without owning a goal. 2 · 22 16 · 176
Age deliberately	Give every memory type an expected lifespan when you design it. 10 · 120

Agent	Owns a goal, chooses its own route, and decides for itself when the work is sufficiently done. 2 · 22 16 · 169
Agent loop	Plan, act, observe, revise. The execution model behind anything that acts rather than answers. 7 · 87 11 · 129 13 · 147 16 · 176
Agentic RAG	The model deciding what to retrieve, iteratively. 9 · 107
Always attach a reason	An unexplained recommendation cannot be evaluated by the person receiving it. 11 · 130
Always render created state	Conversation that writes invisible settings is worse than a settings panel. 8 · 97
Anchor to stable identifiers	Character offsets break under concurrent edits. 6 · 77
Apply optimistically, undo cheaply	The edit lands instantly and reverses instantly. 6 · 77
Ask before persisting anything sensitive	Consent at the moment of capture, not buried in a policy. 10 · 120
ASR	Speech to text. The entry point for voice. 15 · 166
Audit logging	Immutable record of action and authority. 8 · 96 16 · 178
Autonomy levels	A graded scale from suggesting to acting silently. 16 · 176

Autonomy spectrum Suggest, recommend, act on confirmation, act and report, act silently. Autonomy is graded, not binary. 16 · 169

B

Barge-in Letting the user interrupt playback. 15 · 166

Budget the window per turn Decide explicitly what share goes to instructions, memory, retrieval, history. 9 · 109

Build the eval set first Before the feature. It defines what good means. 14 · 157

C

Calibrate the judge Check model grading against human grading on a sample, regularly. 14 · 157

Calibration Aligning stated confidence with actual accuracy. 11 · 129

Cap interruption frequency Recommendation fatigue erodes trust faster than being occasionally wrong. 11 · 130

Capability What a product can do. A domain question, distinct from how the work is packaged architecturally. 2 · 22 17 · 182

Chain of thought Intermediate reasoning steps. 13 · 147

Checkpointing Persisting intermediate state. Lets a long task survive interruption and resume visibly. 13 · 147

Checkpointing and rollback Saved states an action can be reversed to. 16 · 176

Choose the smallest model that passes	Capability you do not need is cost and latency you do pay. 14 · 157
Chunking strategy	How documents are split before indexing. 9 · 107
Citation and span attribution	Linking each claim to its origin. Converts a summary from assertion into something verifiable. 12 · 138
Cite everything	Every claim traceable to a span. No exceptions in high-stakes documents. 12 · 140
Cognitive load	The thinking a product asks of its user. Reducing it is the first theme of the philosophy. 4 · 46
Cognitive unit (CU)	A prompt with a defined input, a specified model, and an expected output. 14 · 156
Complementary strengths	Machines bring volume, patterns and recall; humans bring nuance, values and responsibility. 16 · 169
Component stack	Model, memory, retrieval, tools, skills, agents, evaluation — the layers a product is assembled from. 2 · 22
Compounding value	Memory makes each session worth more than the last, where a stateless product repeats itself. 10 · 111
Confidence estimation	Assessing how sure the system is. 16 · 178
Confidence-determined behaviour	High confidence recommends, medium offers alternatives, low asks, very low stays silent. 11 · 123
Confirm before persisting	The moment a request becomes configuration should be explicit. 8 · 97

Conflict resolution	Handling new facts contradicting stored ones. 10 · 118
Constrain the vocabulary for highstakes actions	Know where to trade expressiveness for reliability. 15 · 168
Constrained decoding	Restricting the token space during generation. 5 · 66 6 · 76 15 · 166
Context degradation	Reduced attention to material mid-window. 9 · 107
Context engineering	Deciding what enters the prompt, in what order. 9 · 107
Context half-life	Every kind of context decays at its own rate, from minutes for what is on screen to years for identity. 9 · 101
Context is a Feature	Principle: context is a designed, first-class product capability, not an implementation detail. 9 · 101
Context window	The attention budget. Determines whether a document can be summarised in one pass or must be chunked. 6 · 76 9 · 107 12 · 138
Conversation over Configuration	Principle: replace settings panels with dialogue, while keeping governance explicit. 8 · 89
Co-pilot the Decision	Principle: the product and the person decide together, with authority placed deliberately. 16 · 169
Cost per interaction	Unit economics of inference. Intelligence carries marginal cost interface work never did. 14 · 157

D

Data residency and erasure	Where personal data lives and how it is removed. 10 · 120
Delegated control	The mark of real agency: the product decides what to attempt next in light of what happened. 2 · 22 16 · 169
Design the repair path first	Misheard input is the normal case, not the edge case. 15 · 168
Deterministic execution	Running settled rules as code, not inference. 8 · 96
Deterministic software	The rules era. Same input, same output; correctness is checked by tests and assertions. 2 · 22
Diff before applying	Machine-check that untouched content is untouched. 6 · 77
Distillation	Training a small model on a large one's output. 14 · 157
Document parsing	Layout-aware handling of tables, clauses, scans. 6 · 76 12 · 138
Drift detection	Monitoring for behaviour change over time. 14 · 156
Dry run	Producing the plan and its effects without committing. 16 · 176
Durable execution	Long-running work outliving a session. 13 · 147

E

Embedding models	Models mapping text to vectors capturing meaning. 9 · 107 10 · 118
Entailment checking	Automated verification that a claim follows from a passage. 12 · 138
Entity resolution	Recognising that references point to the same thing. 9 · 107
Episodic memory	Records of specific past events and interactions. 10 · 118
Eval frameworks	Tooling for faithfulness, relevance, coverage. 12 · 140 14 · 156
Eval sets	Curated cases with known good outcomes. 14 · 156
Explanation faithfulness	Whether a stated rationale is the actual cause. 13 · 147
Extract fields for known types	When you know the document type, extract structure instead of writing prose. 12 · 140
Extractive summarization	Selecting sentences from the source. 12 · 138

F

Filter permissions in the query	Never retrieve then redact. 9 · 109
Fine-tuning	Training on task-specific examples. 14 · 156

Function calling Typed calls the model can make. Settings become tools the system can set on the user's behalf. 7 · 87 8 · 96 15 · 166

G

Gate by consequence and reversibility Not by capability. What the system can do is the wrong question. 16 · 178

Generate several, show one Produce candidates, rank them, surface the best. 11 · 130

Grade outcomes, not wording Success is the result the user wanted, however the system got there. 7 · 88

GraphRAG Retrieval traversing a graph rather than a flat index. 9 · 107

Groundedness Whether claims are supported by the source. 12 · 138

Groundedness evaluation Checking output against source. 6 · 76

Grounding and citation Tying output to retrieved spans. Turns retrieved context into something a user can verify. 9 · 107

Guard models Small classifiers checking safety or policy. 11 · 129

Guardrails Policy enforced outside the model. 7 · 87 16 · 176

H

Hierarchical summarization Summarising in passes, then summarising summaries. 12 · 138

Human in the loop	Designed points requiring human authorisation. 16 · 176
Hybrid search	Combining lexical and semantic scoring. 9 · 107
In-context learning	The model adapting from examples in the prompt. 5 · 66
Index freshness	How current the retrievable corpus is. 9 · 107
Indirect prompt injection	Injection arriving via retrieved documents, web pages, or email. 16 · 178
Infer defaults aggressively	Ask only for what genuinely cannot be derived from context. 7 · 88
Information Age	The preceding era, whose software executed instructions while the human supplied the thinking. 1 · 14 17 · 182
Instrument accept and reject	Without this signal there is no path to improvement. 11 · 130
Intelligence Age	The era in which intelligence, not information, is the abundant resource products are built on. 1 · 14
Intelligence ladder	Transcribe, answer, explain, recommend, act — rising levels of usefulness, not just capability. 11 · 123
Intelligence over Aesthetics	Principle: polish is table stakes; how well a product thinks is the differentiator. 14 · 150
Intelligence product leader	The emerging role that owns intelligence behaviour, not only features and interfaces. 4 · 46

Intent classification	Routing an utterance to a task type. 7 · 87
Intent over Instructions	Principle: let users state the outcome and let the product infer the procedure. 7 · 79
Interface surface area	How much UI a user must operate. As products get more capable, the interface gets quieter. 18 · 189
Interruptibility	The ability to cancel mid-flight. Visible work implies stoppable work. 13 · 149

K

Keep a non-language alternative	Language should be the primary interface, not the only one. 15 · 168
Keep governance out of dialogue	Privacy, security, compliance and policy stay as explicit controls. 8 · 97
Knowledge graph	Entities and relationships as explicit structure. 9 · 107

L

Language over Layout	Principle: natural language becomes the primary interface, and layout the specialist one. 15 · 160
Language takes, layout keeps	Language wins where understanding is needed; direct manipulation wins where doing is enough. 15 · 160
Layer the disclosure	Headline, summary, key clauses, full source, letting the reader stop where they choose. 12 · 140

Least privilege and scoped credentials	Granting only the access a task requires. 16 · 176
Lexical search (BM25)	Classic keyword ranking. Still better than semantic search for names, codes, and exact identifiers. 9 · 107
LLM as judge	Model-based grading at scale. Makes evaluation affordable. 11 · 129 14 · 156
Log authority, not just action	Who authorised it, on what basis, under which policy. 16 · 178
Logprobs	Token-level probabilities. The cheapest available confidence signal, though a rough one. 6 · 76 11 · 129
Long-term memory	State persisting across sessions. What the chapter is actually arguing for. 10 · 118
LoRA and parameter-efficient tuning	Adapting a model via a small trained overlay. 14 · 156
Low-Cognitive First Draft	Principle: begin with thoughtful, structured work instead of a blank page. 5 · 62

M

Make supersession visible	Show what changed when a new instruction overrides an old one. 8 · 97
Make undo cheap before making autonomy wide	Every increment of reversibility buys an increment of trust. 16 · 178

Map confidence to behaviour	High recommends, medium offers alternatives, low asks, very low says nothing. 11 · 130
Mark provenance	Track which changes came from the system so they can be reviewed as a set. 6 · 77
MCP	An open protocol connecting models to tools and data. 7 · 87 9 · 107 16 · 176
Measure effort saved	Track edit distance and discard rate, not accuracy. 5 · 67
Measure retrieval separately	Evaluate whether the right context was found before blaming the model. 9 · 109
Memory conflict resolution	Reconciling instructions contradicting earlier ones. 8 · 96
Memory extraction	Deriving durable facts from raw interaction. 10 · 118
Memory is Part of UX	Principle: what a product remembers, forgets and shows is an experience decision. 10 · 111
Metadata filtering	Constraining retrieval by structured attributes. 9 · 107
Model routing	Directing requests to different models by need. 5 · 67 14 · 157
Multilingual models	Coverage across languages, including mixed input. 15 · 166
Multimodal models	Text, image, and audio in one system. 15 · 166

N

Name the step and count the units Real status requires instrumentation, not animation. 13 · 149

O

Observability Tracing, logging, inspection of live behaviour. 14 · 157

Offer cancel wherever you show progress Progress without control is frustration. 13 · 149

Offline vs online evaluation Pre-release testing versus live measurement. 14 · 156

One clarifying question, maximum Clarification is an exception path. 7 · 88

P

Paint the skeleton first Render structure immediately, fill sections as they stream. 5 · 67

Path stabilisation Naming a route an agent keeps rediscovering. 2 · 22
16 · 176

Permission-aware retrieval Filtering by entitlement at query time. 9 · 107

Place decisive material at the edges	Beginning and end are attended to more reliably than the middle. 9 · 109
Point & Fix	Principle: let users select what is wrong and describe the change, rather than repeat a workflow. 6 · 69
Precision at k	Share of top suggestions that are good. 11 · 129
Predictive software	The learning era. Statistical relationships predict likely outputs; accuracy is the measure. 2 · 22
Preference data	Signals from what users accept and reject. 11 · 129
Preference store	A structured, queryable record of user choices. 8 · 96
Preference tuning	Training on chosen-versus-rejected pairs. 14 · 157
Preview before executing	The cheapest place to catch a misread intent is before it acts. 7 · 88
Procedural memory	Stored knowledge of how the user prefers things done. 8 · 96 10 · 118
Progressive detail	Headline, summary, key details, full source — the summary becomes the interface. 12 · 132
Promote repetition into rules	When a one-off recurs, offer to make it standing. 8 · 97
Prompt caching	Reusing a stable prompt prefix. Makes large persistent context economically viable turn after turn. 5 · 66 6 · 76 9 · 107
Prompt injection	Instructions hidden in content the model reads. 16 · 176

Prompt template	A parameterised, versioned prompt. 5 · 66 8 · 96
Prompt versioning	Treating prompts as versioned artefacts. 14 · 157

Q

Quantization	Reduced numerical precision. 14 · 157
---------------------	---------------------------------------

R

RAG	Retrieval-augmented generation. The overall pattern this chapter describes in prose. 5 · 66 9 · 107
Reasoning models	Models spending extra compute before answering. 13 · 147
Recommend, Don't Respond	Principle: answering a question is the lowest form of intelligence; propose the next step. 11 · 123
Regenerate with a nudge	Offer a directional retry rather than a blank one. 5 · 67
Regression testing	Re-running evals on every change. Model and prompt changes alter behaviour silently. 14 · 156
Rehearse in a sandbox	Assess judgment before granting reach. 16 · 178
Reranking	Ordering by importance. What surfaces first is a ranking problem, not a writing problem. 9 · 107 11 · 129 12 · 138
Re-run on every model change	Including provider updates you did not initiate. 14 · 157

Reserve full reasoning for consequential output	Match depth of explanation to stakes of the decision. 13 · 149
Resolve every utterance to a typed action	Understanding without execution is a demo. 15 · 168
Retrieve wide, rerank hard	Fetch generously, then let a reranker decide what enters the window. 9 · 109
Return patches, not documents	Targeted changes are reviewable, revertible, cheap to verify. 6 · 77
Round-trip latency budget	Total time from speech to response. 15 · 166
Route before you reason	Classify the intent, then choose the handling path. 7 · 88

S

Salience	What counts as important, and to whom. 12 · 140
Salience and write policy	Deciding what merits storage at all. 10 · 118
Sampling temperature	Randomness in token selection. 5 · 67
Sandboxing	Executing against a safe replica. Lets an inferred plan be rehearsed before it touches anything real. 7 · 87 16 · 176

Seed from the user's material	Ground the draft in their most recent or relevant artefact before generating anything. 5 · 62
Segmentation models	Turning a click or box into a precise region. 6 · 76
Self-consistency	Sampling repeatedly and measuring agreement. 11 · 129
Semantic caching	Matching new requests to previously answered similar ones. 13 · 149 15 · 168
Semantic memory	Durable facts about the user or their world. 10 · 118
Semantic search	Retrieval by meaning rather than wording. 9 · 107
Send the neighbourhood	Include surrounding text, position, and structural role alongside the selection. 6 · 77
Separate stable from volatile	A seat preference and next week's itinerary need different treatment. 10 · 120
Session vs persistent state	What lasts a turn, a task, or forever. 8 · 96
Set a hard ceiling before first paint	Decide the maximum acceptable silence and engineer to it. 5 · 67
Ship it deliberately incomplete	Aim for relevant, structured, unfinished. 5 · 67
Short-term memory	State held within a session. Cheap, bounded, discarded. 10 · 118
Show doing, not thinking	"Reviewing 40 documents" beats a rendered monologue. 13 · 149

Show the memory	Make what is remembered visible and correctable in the product. 10 · 120
Skill	A reusable, bounded procedure combining cognitive units, code and tools to perform a known capability. 2 · 22 16 · 176
Slot filling	Inferring missing parameters from available context. 7 · 87
Small language models	Compact models, far cheaper and faster. 5 · 66
Specification gaming	Satisfying the stated goal while missing the intent. 16 · 178
Speculative decoding	A small model proposes, a large one verifies. 5 · 67
Speech-to-speech models	Audio in, audio out, without a text hop. 15 · 166
Streaming	Token-by-token delivery. The simplest and most effective form of visible progress. 5 · 66 13 · 147
Structured extraction	Pulling named fields rather than prose. 12 · 138
Structured output	Schema-constrained generation. Intent resolves to typed parameters, not a paraphrase. 5 · 66 6 · 76 7 · 87 8 · 96
Summarise reasoning, never dump it	Condense the trace into something a person would say. 13 · 149

Summarization compaction

Condensing history to fit a budget. 10 · 118

T**Task decomposition**

Breaking a goal into ordered sub-goals. 7 · 87

Teach affordances by example

Users cannot see what they may say. 15 · 168

Test for knowing when to stop

The hardest capability to build, the easiest to omit from evaluation. 16 · 178

The delight shift

Delight moves from aesthetic polish to how intelligently a product understands and acts. 14 · 150

The Devil is in the TL;DR

Principle: summarisation is where trust is won or lost, so make omission accountable. 12 · 132

Thinking budget

How much reasoning to reveal, set by the stakes of the outcome and the length of the task. 13 · 141

Thoughtware

Software that understands, adapts and reasons toward goals rather than executing instructions. 2 · 22

Three eras of software

Deterministic, predictive, thoughtware — a taxonomy of how software decides, fails and is tested. 2 · 22

Three gates

System confidence, consequence of failure, reversibility — together they place an action on the spectrum. 16 · 169

Three responsibilities

A cognitive unit decides, a skill performs, an agent pursues. Responsibilities, not sizes. 2 · 22

Time to first token	Delay before output starts. The summary must appear before the user gives up and scrolls. 5 · 66 12 · 140 13 · 147
Tokens per second	Sustained output rate. Determines whether visible progress feels like progress or like struggle. 5 · 66 13 · 147
Tool schemas	Typed declarations of available capability. 7 · 87
Tool use	The model's ability to take real action. 11 · 129 16 · 176
Tool-call traces	The record of actions taken. 13 · 147
Tracing and spans	Structured instrumentation of execution. 13 · 147
Track cost per successful outcome	Not cost per call. Failures are expensive twice. 14 · 157
Trajectory evaluation	Grading the path an agent took, not just its text. 7 · 88 16 · 178
TTL and decay	Ageing and expiry rules for stored items. 10 · 118
TTS	Text to speech. Output quality shapes trust more than most teams anticipate. 15 · 166

U

User-editable memory	Interfaces to view, correct, and delete. 10 · 120
-----------------------------	---

V

Vary salience by role	A lawyer and a buyer need different summaries of the same contract. 12 · 140
------------------------------	--

Vector database	Storage and nearest-neighbour search over embeddings. 9 · 107
Vector store vs profile store	Fuzzy recall versus structured facts. 10 · 118
Verify entailment automatically	Check faithfulness before display, not in review. 12 · 140
Visible Thinking	Principle: reveal enough reasoning to be trusted, without narrating every thought. 13 · 141
Vision-language models	Models that read images. “Fix this bit” requires the system to see the bit. 6 · 76 12 · 138 15 · 166
Voice activity detection and endpointing	Knowing when a person started and stopped speaking. 15 · 166

W

Word error rate	Standard ASR accuracy metric. Useful, but secondary. 15 · 168
Write less, recall better	Storage is easy. Surfacing the right memory at the right moment is the hard problem. 10 · 120

This is the end

Thank you