



Your Data and AI

(AIGuide 12)

Understanding What We Give AI—and What May Follow

AI Sourced Facts (AISF) Pte. Ltd.

© AI Sourced Facts (AISF)

All rights reserved.

Edition: Version 1.0

Publication Date: August 2026

Preface

When people interact with artificial intelligence, attention naturally falls on what the AI provides.

A question produces an answer. A document produces a summary. An image produces an interpretation. A request produces assistance.

Yet every such interaction has another dimension.

Information moves in the other direction too.

People type words, describe circumstances, upload documents, submit images, provide context and introduce information about themselves, their work and sometimes other people. The information supplied to accomplish a task may be broader than the task itself.

Some information is obvious.

Other information can be revealed through language, context and patterns of interaction. Information may support inferences that were never explicitly stated. It may have significance beyond the immediate exchange. Separate pieces of information can combine into a larger picture. Increasing personalisation can make information more relevant to continuing interaction.

And what appears to the individual as a simple conversation may therefore form part of a wider information relationship.

That relationship is not inherently beneficial or harmful.

Information enables many of the capabilities that make AI useful. Context can improve relevance. Documents and other materials allow AI to work directly with complex information. Continuity can reduce repetition and support more sustained assistance.

But usefulness does not remove the need for understanding.

A person can successfully operate an AI system without necessarily seeing the complete informational environment surrounding that interaction.

This Guide examines that difference.

Its purpose is not to tell readers what information they should or should not provide.

It is to help make the information relationship more visible.

Purpose of the Guide

The purpose of **AIGuide 12 — Your Data and AI** is to help readers understand the information dimension of interacting with artificial intelligence.

The Guide begins with a simple proposition:

Every AI interaction has an information dimension.

From that foundation, it examines:

what people deliberately give AI;

information that may be revealed without being deliberately provided;

the distinction between information and inference;

what it means when information persists beyond an immediate interaction;

personalisation and the accumulation of informational context;

the difference between conversation and confidentiality;

consent within complex information environments;

the visibility gap between what users can readily observe and the wider information relationship;

how small pieces of information can combine into a larger picture;

information about people other than the person interacting with AI;

what it means to say that AI “knows” something about a person;

the growing dependence of richer AI interaction upon information; and

how people can live with an information relationship they may never see completely.

This is not a privacy manual, cybersecurity guide or legal guide.

It does not prescribe what information individuals should provide to particular AI systems, recommend settings or providers, or attempt to establish universal behavioural rules.

Different systems, technical architectures, organisational environments, jurisdictions and purposes can involve materially different considerations.

The concern of AIGuide 12 is more foundational:

When we interact with AI, what information relationship are we creating—and how well do we understand it?

Global Neutrality Statement

Artificial intelligence does not operate within a single global information environment.

Laws, regulation, cultural expectations, institutional practices, technical infrastructure and attitudes towards privacy and personal information differ substantially across jurisdictions and communities.

AI systems also differ.

Information handling, persistence, personalisation, contextual capabilities and operating arrangements may vary between providers, products, environments and versions. They may also change over time.

AIGuide 12 therefore does not assume that information supplied to one AI system is necessarily treated in the same way as information supplied to another.

Nor does it attempt to convert jurisdiction-specific legal concepts into universal principles.

Terms such as *data*, *privacy*, *consent*, *confidentiality*, *memory* and *personalisation* can have technical, legal and everyday meanings that are not always identical.

The Guide uses these concepts to improve understanding rather than to provide jurisdiction-specific legal conclusions.

Its emphasis remains on the underlying human question:

What does a person need to understand about the informational character of AI interaction in order for human judgement to remain meaningful?

Editorial Standard

AISF AI Guidance is designed to be **structured, analytical, neutral and non-prescriptive**.

AIGuide 12 therefore avoids two opposite assumptions.

It does not assume that giving information to AI is inherently harmful.

Nor does it assume that greater information availability and personalisation are automatically desirable.

Information enables capability.

It can also create informational complexity.

The Guide distinguishes between:

information and consequence;

deliberate provision and indirect revelation;

information and inference;

persistence and human memory;

conversation and confidentiality;

participation and complete understanding;

visible interaction and wider informational context;

information quantity and human understanding; and

apparent knowledge and computational use of available information.

These distinctions allow the subject to be examined without converting awareness into prescription.

Understanding what may be occurring does not itself determine what a person should do.

It improves the basis upon which judgement can be exercised.

How to Use This Guide

AIGuide 12 is designed as a progressive examination of the information relationship between people and AI.

The opening chapters establish the information that enters an AI interaction directly and indirectly.

The Guide then moves from information to **inference**, and from the immediate interaction to questions of **persistence and personalisation**.

The middle chapters examine concepts that can easily be misunderstood when conversational interfaces resemble familiar human interaction: confidentiality, consent and visibility.

Later chapters broaden the picture by examining aggregation, information concerning other people and the everyday language of AI “knowing” a person.

The final chapters bring these elements together by examining increasing information dependence and the challenge of living with an informational relationship that may remain only partly visible.

Readers may consult chapters independently.

However, reading the Guide progressively reveals an important development:

what begins as a visible exchange of information can become a richer, more contextual and more complex information relationship.

Table of Contents

Preface

Purpose of the Guide

Global Neutrality Statement

Editorial Standard

How to Use This Guide

Version & Governance

Chapters

Chapter 1 — Every AI Interaction Has an Information Dimension

Chapter 2 — What We Give AI

Chapter 3 — Information We Reveal Without Deliberately Providing It

Chapter 4 — From Information to Inference

Chapter 5 — When Information Persists

Chapter 6 — Personalisation and the Accumulating Picture

Chapter 7 — The Difference Between Conversation and Confidentiality

Chapter 8 — Consent in a Complex Information Environment

Chapter 9 — The Visibility Gap

Chapter 10 — When Small Pieces Become a Larger Picture

Chapter 11 — Other People's Information

Chapter 12 — What Does AI Know About You?

Chapter 13 — The Growing Data Dependence of AI Interaction

Chapter 14 — Living with an Information Relationship We Cannot Fully See

Key Ideas from This Guide

AI Guidance — Continuing the Exploration

Future Edition Updates & User Submissions

Acknowledgements

AISF Universal Disclaimer

About AI Sourced Facts (AISF)

Publication & Governance Information

Version & Governance

AIGuide 12 — Your Data and AI

Understanding What We Give AI—and What May Follow

Edition: 2026 Edition

Series: AISF AI Guidance

Published by: AI Sourced Facts (AISF)

This Guide forms part of the continuing AISF AI Guidance series.

Artificial intelligence systems, information practices, contextual capabilities, personalisation features and regulatory environments will continue to evolve.

Future editions may therefore reflect material developments where they affect the continuing relevance or clarity of the subjects examined.

Such revisions do not alter the underlying editorial principles of the series: neutrality, non-endorsement, structured analysis, global applicability where appropriate, human-centred interpretation and recognition of uncertainty.

Editorial oversight of the AI Guidance series is maintained by AISF.

Chapter 1 — Every AI Interaction Has an Information Dimension

For many people, interacting with artificial intelligence feels straightforward. A question is typed into a conversational interface, a document is uploaded for examination, an image is submitted for interpretation, or a request is made for assistance with a particular task. The AI responds, and the interaction may appear to consist simply of an exchange: the person provides something, the system produces something in return.

Yet every such interaction also has an information dimension.

This dimension is easy to overlook because attention naturally falls on what the AI produces. A person asking for an explanation is interested in the explanation. Someone requesting a summary is concerned with the summary. A user seeking help drafting a message is likely to focus on whether the resulting text meets the intended purpose. The information supplied to make that interaction possible can seem secondary to the outcome being sought.

Understanding AI interaction more fully requires looking at both sides of that exchange. There is the visible output produced by the system, but there is also the information introduced by the person and the informational context within which the interaction occurs. What is entered, uploaded, revealed or otherwise made available can form part of an information relationship between the individual and the AI system.

That relationship is the subject of this Guide.

More Than Asking a Question

A prompt can look like nothing more than an instruction. “Explain this concept.” “Summarise this report.” “Help me prepare for a meeting.” “Translate this message.” But prompts frequently contain information as well as instructions.

Consider the difference between asking an AI to “draft a general email requesting a meeting” and asking it to “draft an email to my manager explaining that I need to move tomorrow’s performance review because of a family matter.” Both are requests for writing assistance, but the second interaction contains considerably more information about the person and their circumstances.

The same distinction arises across many forms of AI use. A request for help comparing two generic options differs informationally from a request that includes a person’s income, location and financial circumstances. Asking an AI to explain a medical term is different from uploading a medical report containing a name, date of birth and clinical history. Requesting a summary of an invented document differs from uploading an internal workplace document containing information about colleagues, customers or commercial activities.

These differences do not mean that one form of interaction is necessarily appropriate and another inappropriate. That would turn an examination of the information relationship into a set of behavioural recommendations, which is not the purpose of this Guide. The important

observation is more fundamental: **what appears to the user primarily as a task may simultaneously be an exchange involving information.**

Recognising that dual character is the starting point for understanding data and AI.

Information Can Take Many Forms

The word *data* can sound technical. It may suggest databases, spreadsheets, computer records or large collections of numbers. In everyday AI interaction, however, the relevant concept is considerably broader.

A sentence can contain information. So can a photograph. A voice recording may contain information not only in the words spoken but also in characteristics associated with the recording. A document can contain names, dates, relationships and circumstances. A series of questions can reveal interests or concerns. An instruction can contain details about a workplace, family situation, project or personal objective.

Information therefore does not need to look like conventional “data” to have an informational character.

This matters particularly with conversational AI because natural language makes interaction feel familiar. People are accustomed to communicating through words, and conversation is one of the principal ways human beings exchange information with one another. When the same familiar form is used to interact with an AI system, the informational character of the exchange may be less noticeable than it would be if a person were completing a formal database record or entering information into clearly labelled fields.

The interface may feel like conversation. The interaction nevertheless involves information.

This does not mean that a conversational AI should simply be understood as a database, nor that every piece of information introduced into every AI system is handled in the same way. AI systems, providers, technical architectures and operating environments differ. Information practices can also change over time and across jurisdictions, products and settings. The purpose at this stage is not to generalise about what any particular system does with information. It is to recognise that information exists within the interaction in the first place.

The Information Relationship Is Wider Than the Prompt

It is also tempting to equate the information relationship with whatever a person consciously types into a prompt. That is the most visible component, but it is not necessarily the complete conceptual picture.

People may provide information through uploaded documents, photographs, audio, video or other materials. Information may arise from the context of an interaction. Patterns of engagement can themselves have an informational character. Systems may, depending on their design and operation, be capable of drawing inferences from information rather than relying only on facts explicitly stated by the user.

These dimensions will require careful examination later in the Guide. At this point, their importance lies in demonstrating why the simple mental model of “I typed a question and received an answer” may be incomplete.

A more useful conceptual starting point is to recognise an **information relationship**.

The term does not imply that an AI system has a human relationship with the user. Nor does it suggest that AI “knows” a person in the ordinary human meaning of knowing someone. Such language can easily anthropomorphise computational systems and obscure important distinctions.

Instead, an information relationship describes the fact that interaction can involve information concerning the person, their circumstances, their activities, other people or the context in which the interaction occurs. The nature and extent of that relationship can vary considerably.

A person making a single generic request may create a very limited informational context. Someone repeatedly using an AI-enabled service in a more personalised environment may experience something substantially richer. Between these points lies a wide range of possible interactions.

Why the Relationship Can Be Difficult to Notice

The information dimension of AI interaction can remain partly invisible because usefulness tends to dominate attention.

When a system provides an immediate and relevant response, the user's attention is understandably directed towards the benefit received. The process can become routine. As familiarity increases, the act of providing information may become embedded within the task rather than experienced as a separate event.

This is not unique to AI. Many digital services depend on information while presenting themselves primarily through the functions they perform. AI nevertheless adds particular complexity because interaction can be open-ended. Rather than selecting only from predetermined fields, users can communicate in ordinary language, upload rich material and provide extensive context.

The flexibility that makes AI useful can therefore also broaden the informational character of the interaction.

There is another complication. Different participants may see different parts of the information environment. The individual sees the interface and the material they deliberately provide. The underlying system may involve technical processes, policies and arrangements that are not equally visible from that interface. Ordinary users cannot necessarily be expected to understand those processes merely because they can operate the system successfully.

Being able to use an AI system and understanding its information relationship are therefore not the same thing.

Information Does Not Mean Consequence

Recognising an information relationship should not lead automatically to conclusions about harm, misuse or risk.

Information can enable useful functions. Context can make responses more relevant. Personalisation may depend upon information. Documents must ordinarily be made available in some form if an AI system is expected to analyse their contents. Many valuable AI interactions would be difficult or impossible without information.

At the same time, the existence of useful purposes does not eliminate questions about what information is involved, what can be derived from it, whether it persists, how it may contribute to personalisation, or how clearly the individual understands the surrounding information environment.

Those questions should not be collapsed into a simple judgement that information use is either “good” or “bad.” Different circumstances, systems, purposes and individuals can produce different considerations.

This Guide therefore treats the information relationship as something to be understood before conclusions are drawn about it. That approach reflects a broader distinction between awareness and prescription. Understanding what may be occurring does not itself determine what an individual should do.

A Relationship That May Develop Over Time

The informational significance of an interaction may also be difficult to understand when each exchange is considered in isolation.

A single question might reveal very little. A single uploaded document might have a clearly defined purpose. One detail about an individual's preferences may appear insignificant. Yet information can acquire different meaning when considered alongside other information.

This introduces questions that extend beyond the immediate interaction. Can separate pieces of information contribute to a larger picture? What is the significance of persistence? How does personalisation relate to accumulated context? What can be inferred rather than directly stated? What happens when information supplied by one person concerns somebody else?

These are not questions that can be resolved simply by examining the wording of one prompt.

They require an understanding of the broader information relationship.

The distinction becomes increasingly important as AI interaction moves beyond occasional generic questions towards systems capable of providing more contextual and personalised assistance. Greater contextual awareness can make AI interaction more useful, but it also makes the informational foundations of that interaction more significant.

Neither consequence should be assumed to outweigh the other. They are different dimensions of the same developing relationship.

Making the Invisible More Visible

The central purpose of examining data and AI is therefore not to encourage suspicion of AI systems or confidence in them. It is not to establish rules about what people should disclose, identify particular systems that should be used, or provide legal, privacy or cybersecurity advice. Those purposes fall outside the boundaries of this Guide.

The more basic task is to make an often-overlooked dimension of AI interaction easier to see.

When a person interacts with AI, the visible exchange may be a question and an answer, a document and a summary, an image and an interpretation, or an instruction and a generated result. Beneath that functional description lies another question:

What information is becoming part of this interaction?

Once that question is recognised, others naturally follow. Some information is deliberately provided. Some may be revealed indirectly. Some may concern other people. Some may contribute to inference or personalisation. Some may persist beyond the moment in which it was introduced. And the individual may not always have a complete view of the wider informational context surrounding the interaction.

Those later questions require separate examination. The foundation, however, comes first.

Every AI interaction has an information dimension.

Understanding that simple proposition changes the way the subject can be examined. The interaction is no longer viewed only in terms of what AI gives the person. It can also be viewed in terms of what information the person brings into the interaction, what additional information may emerge from it, and what kind of informational relationship may consequently develop.

That does not answer whether the relationship should be welcomed, limited, trusted or changed.

It establishes something more fundamental: **before people can form informed judgements about their relationship with AI, they need to be able to see the information dimension of that relationship clearly enough to understand that it exists.**

Chapter 2 — What We Give AI

The information relationship between a person and an AI system begins most visibly with what the person provides.

Sometimes this is obvious. A user types a name, describes a situation, uploads a document or submits a photograph. At other times, what is being provided may seem less like “data” and more like ordinary conversation. A person explains a problem, supplies background information, pastes an email, describes a workplace situation or asks for help thinking through an issue.

In each case, information becomes part of the interaction.

Chapter 1 established that every AI interaction has an information dimension. The next step is to examine the most readily observable part of that dimension: information that people deliberately provide when seeking an AI system's assistance.

This information can take many forms. It can concern the user directly, another person, an organisation, a situation, a document or an event. It can be highly personal or entirely impersonal. It may consist of a few words or an extensive collection of materials. And it may be provided because the information itself is the subject of the task or simply because additional context appears likely to produce a more useful response.

Understanding these distinctions helps reveal how much informational variety can sit behind the apparently simple act of “using AI.”

The Prompt Is Also Information

A prompt is usually understood in functional terms. It tells an AI system what the user wants.

“Explain this.”

“Compare these alternatives.”

“Summarise this document.”

“Help me understand why this happened.”

Yet a prompt can perform two functions simultaneously. It can instruct the system, and it can provide information.

A generic request such as “Explain compound interest” may communicate relatively little about the person asking it. A request such as “I am approaching retirement and have these three investments; explain how compound interest has affected each of them” contains considerably more context.

The purpose of the second prompt may still be explanation, but the interaction now includes information about the individual's circumstances.

This distinction becomes even clearer when prompts are conversational. People often improve an AI response by adding context: what they are trying to accomplish, what has already happened, who is involved, what constraints exist and why the matter is important.

The resulting interaction may become richer and more useful because the system has more information with which to work.

But informational richness and functional usefulness are different concepts. Recognising that more context has been supplied does not determine whether providing it was appropriate. Nor does it establish what any particular system will subsequently do with it. It simply makes visible something that can otherwise disappear into the familiarity of conversation: **context is information too.**

Personal Information Is Often Embedded in Context

People do not necessarily provide personal information by completing a field labelled “personal information.”

It can appear naturally within a request.

Someone might mention their age when discussing education or retirement, their occupation when asking about a workplace issue, their country when seeking an explanation of a local situation, or family circumstances when asking for help organising an event. A person might include a name because it makes a draft letter easier to prepare, or describe previous experiences because they seem relevant to the response sought.

Much of this can feel incidental.

That is significant because people may have a relatively narrow mental model of personal information. Formal identifiers such as names, addresses, identification numbers and contact details are readily recognised as information about a person. Yet a description of someone's circumstances can also be informative even when those obvious identifiers are absent.

A combination such as occupation, approximate age, location, family circumstances and a description of a particular event may say considerably more about an individual than any one detail considered alone.

The significance of combining information will be examined later in this Guide. At this stage, the important point is that information does not stop being personal merely because it appears inside an ordinary sentence rather than a formal record.

Natural-language AI interaction makes this especially relevant. People can explain themselves freely rather than communicating only through predetermined categories. That flexibility can make an interaction easier and more expressive. It can also mean that personal information becomes woven into the conversation without being experienced as a separate act of data provision.

Documents Bring Their Own Information

AI interaction increasingly extends beyond typed prompts.

A user may upload a report for summarisation, a contract for explanation, meeting notes for organisation, a spreadsheet for analysis, correspondence for drafting assistance, or a résumé for review. In such cases, the information supplied to the AI system is not limited to the instruction accompanying the upload.

The document itself becomes part of the informational context.

That distinction matters because documents often contain substantially more information than is necessary to describe the task. A report may identify its authors and contributors. Meeting minutes can name participants and record what they said. Correspondence can contain addresses, signatures and previous messages. A résumé may include employment history, qualifications, locations and contact information. Workplace material can contain information about customers, employees, suppliers or internal activities.

The person uploading the document may be interested in only one aspect of it. The material nevertheless contains the other information as well.

Again, this observation does not establish what happens to that information after it is supplied. Different systems and environments may handle uploaded materials differently. Nor does it imply that documents should or should not be uploaded.

The narrower conceptual point is that **the informational scope of an interaction can be wider than the user's immediate purpose.**

A person may think, “I provided a report to be summarised.” From an information perspective, the interaction also includes whatever information was contained within that report.

Images, Audio and Other Rich Material

The same principle applies beyond documents.

An image submitted to an AI system may contain information about people, places, objects, surroundings or events. A screenshot can contain text, account details, notifications or portions of other applications that were not the principal reason for sharing it. A photograph taken for one purpose can contain background information that is easy for the person submitting it to overlook.

Audio can similarly carry more than the words a user intends an AI system to process. Depending on the material and the capabilities involved, recordings can contain voices, background sounds, conversations and contextual details.

Video can combine many of these informational dimensions.

These richer forms of input illustrate an important feature of contemporary AI interaction: information can be supplied without first being converted into a conventional text description.

That makes AI systems more capable of working with the kinds of materials people encounter in everyday life. It also broadens what “giving information to AI” can mean.

A prompt may consist of ten words accompanying an image containing far more informational detail than those ten words convey. A simple request to “summarise this” may accompany a document running to dozens of pages. The visible instruction therefore cannot always be treated as a measure of the amount or variety of information introduced into an interaction.

Information About Other People

What people give AI is not necessarily information about themselves.

A manager asking for help organising meeting notes may provide information about employees. A person seeking help composing a message may describe a disagreement with a friend. Someone summarising correspondence may include information written by another person. A family member may describe the circumstances of relatives. A consultant may work with material concerning clients.

This introduces an important distinction between **the person interacting with the AI system and the people represented in the information supplied**.

They are not always the same.

In ordinary conversation, people routinely discuss others. Documents routinely contain multiple names and perspectives. Photographs can depict several people. Workplace materials often concern individuals other than whoever happens to upload them.

AI interaction can therefore involve information relationships extending beyond the immediate user.

The broader implications of third-party information require dedicated examination later in this Guide. They also intersect with other contexts, including family and organisational relationships, that have their own considerations. At this stage, the purpose is simply to establish that the phrase “your data and AI” should not be interpreted as meaning only information about the person operating the system.

What **you give AI** may include information about **someone else**.

The Difference Between Purpose and Content

A recurring theme across these examples is the distinction between why information is provided and what information is actually present.

The purpose may be narrow:

summarise a report;

improve a letter;

explain a photograph;

analyse a spreadsheet;

organise meeting notes.

The content supplied to accomplish that purpose may be much broader.

This difference is easy to miss because human attention is goal-directed. When someone uploads a document to obtain a summary, they naturally think of the document as material for summarisation. When they paste an email to improve its wording, they see it primarily as text requiring editing.

An information perspective asks an additional question: **What does the material contain?**

The answer may include information that is central to the task, information that is merely contextual, and information that is incidental to the user's purpose.

These categories are not fixed. A person's name may be essential in one interaction and irrelevant in another. Location may be necessary context for one question and incidental detail for another. The same document may be used for different purposes by different people.

The informational character of AI interaction is therefore highly contextual.

Deliberate Provision Does Not Mean Complete Awareness

The information examined in this chapter is deliberately supplied in the broad sense that the person types, uploads, records or otherwise introduces material into an AI interaction.

But deliberate provision should not be confused with complete awareness of everything the material communicates.

A person may knowingly upload a photograph while paying little attention to its background. Someone may deliberately provide a document without recalling every name or detail contained within it. A user may consciously describe a situation without recognising how much information can be derived from the combination of details provided.

This creates an important transition.

There is a difference between **information we intentionally introduce into an interaction** and **information that may be revealed, interpreted or inferred through that interaction**.

The first is the primary territory of this chapter. The second begins to take the Guide into a more complex informational environment.

Understanding what we give AI therefore starts with the visible: prompts, personal details, documents, images, audio, contextual explanations and other material consciously brought into an interaction.

But it cannot end there.

The act of providing information may be deliberate while the full informational significance of what has been provided remains less obvious. Words can communicate more than their immediate purpose. Documents can contain more than the section being examined. Images can reveal more than their intended subject. Information about one person can carry information about others.

The apparently simple question **“What did I give the AI?”** can therefore have a more complicated answer than **“What did I type?”**

Recognising that difference prepares the ground for the next stage of understanding: information that an interaction may reveal even when the person did not deliberately set out to provide it.

Chapter 3 — Information We Reveal Without Deliberately Providing It

When people think about the information they give an AI system, they are likely to think first about deliberate acts: typing a prompt, entering personal details, uploading a document, submitting an image or providing background information. Chapter 2 examined this visible side of the information relationship.

But deliberate provision is only part of the picture.

Human communication frequently reveals more than people explicitly set out to communicate. The words someone chooses, the subjects they return to, the sequence of their questions, the context they provide and the patterns within repeated interactions can all carry information. In ordinary human relationships, this is familiar. People learn things about one another not only from direct statements but also from patterns of behaviour, language, interests and choices.

AI interaction introduces a related but distinct informational possibility. Depending on the system, its design and the context in which it operates, information may arise from **how a person interacts**, not merely from what that person consciously intends to disclose.

This does not mean that every AI system observes, retains or uses every possible feature of an interaction. Systems differ substantially, and the technical and organisational arrangements surrounding them can also differ. Nor does the existence of information establish that a particular consequence follows from it.

The important conceptual step is narrower: the information relationship can extend beyond what a person deliberately thinks of themselves as providing.

Disclosure Is Not Always a Formal Act

The idea of “disclosing information” can suggest a deliberate event. A person is asked a question and provides an answer. A form requests a date of birth, address or occupation. A user chooses to upload a document.

Conversational interaction is less neatly structured.

Suppose someone repeatedly asks an AI system to explain material in simpler language. The person has not necessarily stated anything about their educational background, confidence or familiarity with the subject. Yet the pattern of requests itself contains information about the way they are engaging with the material.

Another person may repeatedly ask questions about a particular industry, city, hobby or type of problem. No single question necessarily declares, “This subject is important to me.” Taken together, however, the interactions can reveal a pattern of interest.

Similarly, a user may describe an event without explicitly stating how they feel about it, yet the language chosen can carry contextual information. A person may mention a sequence of

workplace difficulties without directly identifying their broader concern. Someone may ask several questions about relocation without ever stating that they are considering moving.

These examples do not establish what an AI system concludes from such interactions. They demonstrate a more basic distinction between **explicit statement** and **informational revelation**.

People can reveal something without formally declaring it.

Behaviour Can Have an Informational Character

Information is often associated with content: the words in a prompt, the details in a document or the objects visible in an image.

Behaviour can also be informative.

In digital environments generally, patterns of interaction can include such things as frequency, sequence, choices and repeated areas of attention. In an AI context, the relevant possibilities may include which subjects a person repeatedly explores, how an interaction develops, what kinds of assistance are requested, which questions are refined and what contextual information is repeatedly introduced.

The significance of these patterns varies. Some may say very little. Others may become more meaningful when considered over time or alongside other information.

This does not require treating every ordinary action as deeply revealing. The point is that behaviour can form part of an informational environment even when the person experiencing that behaviour thinks primarily in terms of performing tasks.

A user may think, “I asked five questions.”

Viewed differently, those five questions also constitute a pattern of interaction.

The distinction matters because the user's mental model may focus on individual exchanges while information can sometimes be understood through relationships among exchanges.

That possibility becomes increasingly relevant as AI experiences become more contextual or personalised. The more an interaction is capable of drawing upon context, the more important it becomes to distinguish information deliberately stated from information represented by patterns within the interaction.

Language Carries More Than Literal Facts

Natural language is particularly rich in information.

People do not communicate only through formal factual statements. Vocabulary, phrasing, references, examples and context can all contribute to meaning. A prompt may reveal familiarity with a subject without saying “I am an expert.” Another may suggest limited familiarity without stating “I am a beginner.”

Language can also contain clues about professional or social context. A person may use specialised terminology associated with a particular field. References to local institutions, practices or events may provide contextual indications about location or environment. Descriptions of routines can reveal aspects of daily life.

None of these clues necessarily produces a reliable conclusion. A person can use terminology from a profession in which they do not work. Someone can ask questions about a country in which they do not live. A user may explore a subject purely out of curiosity.

This uncertainty is important.

Information that can contribute to an interpretation is not the same as information that proves the interpretation correct.

That distinction prevents the concept of indirect revelation from being overstated. AI interaction may contain signals from which something could potentially be interpreted, but those signals can be incomplete, ambiguous or misleading.

The informational relationship therefore includes uncertainty as well as information.

Context Can Reveal Relationships

The context people provide can also reveal relationships between pieces of information.

Consider a user who says:

“I need help preparing for a meeting tomorrow. My manager will be there, along with the head of our regional office. We are discussing the project I have led for the past six months.”

The immediate purpose may simply be meeting preparation. Yet the description also communicates something about the person's organisational relationships, responsibilities and recent activities.

No formal profile was supplied. The user did not necessarily intend to provide an account of their position within an organisation. The details were offered because they made the request easier to understand.

This illustrates why the boundary between deliberate and indirect provision is not absolute.

The user deliberately typed every word. In that sense, the information was consciously introduced. But the person may have intended to communicate one thing—the context needed for useful assistance—while simultaneously revealing other things that were not the focus of the interaction.

The information relationship is therefore not divided neatly into “information knowingly provided” and “information secretly obtained.” There is a broad middle ground in which information is present in material deliberately supplied but its wider significance is not necessarily the user's focus.

Repetition Can Change Informational Meaning

A single piece of information and a repeated pattern do not always carry the same meaning.

One question about starting a business may reflect casual curiosity. A sustained series of questions about business structures, market research, financing and possible locations may provide a different informational picture.

One request for help writing a job application says something about the immediate task. Repeated requests concerning particular roles, interviews and employment conditions can create a broader context.

One question about a travel destination may be general research. A sequence concerning accommodation, schools, neighbourhoods, employment and moving arrangements may suggest something different.

“Suggest” is important here.

Patterns can support interpretations without establishing certainty. There may be alternative explanations for the same behaviour. A person could be researching on behalf of somebody else, writing fiction, conducting professional research, teaching a subject or simply exploring possibilities.

The existence of a pattern therefore does not make an interpretation true.

This distinction also helps preserve the boundary between AIGuide 12 and questions of truth and verification addressed elsewhere in the AIGuide series. The concern here is not how a conclusion should be verified. It is that patterns of interaction can themselves become informationally meaningful, whether or not conclusions drawn from them prove accurate.

The Difference Between Information and Inference

At this point, two concepts begin to approach one another.

The first is **information revealed indirectly**. The second is **information inferred from what is available**.

They are related, but they are not identical.

A pattern of repeated questions is information about the interaction. A conclusion that the person has a particular intention, preference or characteristic would go further. It would be an inference drawn from that information.

Similarly, the use of specialist terminology is observable within the interaction. Concluding that the user works in the relevant profession is an inference.

Maintaining this distinction is important because it prevents observable information from being confused with conclusions derived from it.

Chapter 4 will examine inference directly. For now, the relevant point is that indirect revelation can provide material from which inferences may become possible.

This also explains why a person's answer to the question “What did I tell the AI?” may not capture the full informational character of an interaction.

The person may remember the facts deliberately stated. The interaction may contain additional patterns and contextual signals that were not experienced as separate disclosures.

The User's Perspective and the System's Perspective

One reason this distinction can be difficult to appreciate is that the person and the computational system do not necessarily encounter an interaction in the same conceptual way.

A human user experiences purpose.

“I am trying to understand this.”

“I am preparing for a meeting.”

“I am exploring an idea.”

“I am asking for help with a document.”

The informational environment can be described differently. It contains text, context, uploaded material and potentially other characteristics associated with the interaction, depending on the system involved.

Neither description is inherently wrong. They represent different perspectives on the same event.

The human perspective explains **why the interaction is taking place**.

The information perspective asks **what the interaction contains or reveals**.

Understanding AI and data requires being able to hold both perspectives at once.

A Wider Information Relationship

Recognising indirect revelation does not require assuming that AI interaction is uniquely intrusive or that users should approach every exchange with concern. Human activities have always generated information beyond what individuals consciously intend to communicate, and digital environments have long added further informational dimensions.

What AI changes is the form and richness of some interactions.

People can communicate extensively in natural language. They can refine requests through extended dialogue. They can introduce large amounts of context and combine text with documents, images, audio and other materials. As AI systems become more capable of

contextual interaction, the distinction between the information deliberately supplied for a task and the broader information represented by the interaction can become increasingly important.

The information relationship introduced in Chapter 1 is therefore already wider than the visible prompt.

Chapter 2 examined what people consciously give AI. This chapter adds a second layer: **what interaction itself can reveal**.

The difference is subtle but fundamental.

A person may deliberately provide words without deliberately intending every informational implication those words can carry. They may ask individual questions without thinking of the questions collectively as a pattern. They may provide context for one purpose while that context communicates something broader about their circumstances.

And even then, another step remains.

Information—whether directly provided or indirectly revealed—can become the basis for conclusions that were never stated at all.

That is where the information relationship moves from **revelation** to **inference**.

Chapter 4 — From Information to Inference

People can provide an AI system with information without necessarily providing every conclusion that might later be drawn from it.

A person may state their age without stating that they are approaching a particular life stage. They may describe several workplace events without identifying a pattern among them. They may ask repeated questions about a subject without explaining why that subject matters to them. A document may contain several facts which, when considered together, suggest something that no individual sentence explicitly says.

This is the territory of **inference**.

Chapter 2 examined information people deliberately provide. Chapter 3 extended the picture to information that interaction can reveal without being the person's deliberate focus. A further distinction is now necessary. Information can become the basis for conclusions, classifications, associations or possibilities that were never directly supplied at all.

Understanding this movement from information to inference is important because it changes the apparent boundaries of the information relationship. The question is no longer simply, “What information did I provide?” It also becomes, “What might potentially be derived from the information that is available?”

That does not mean every possible inference is made, retained or acted upon by an AI system. Nor does it mean an inference is necessarily correct. Systems differ, purposes differ and circumstances differ. The important point is conceptual: **information supplied about a person and information inferred about that person are not the same thing.**

Information and Inference Are Different

Suppose someone tells an AI system:

“I have worked for the same organisation for 28 years and am thinking about what I might do next.”

Several pieces of information have been explicitly provided. The person has a long employment history with one organisation and is considering some form of change.

Other conclusions might appear plausible. Perhaps the person is contemplating retirement. Perhaps they are considering another job. Perhaps they are thinking about reducing their workload, starting a business or pursuing an entirely unrelated interest.

None of those possibilities was actually stated.

They are interpretations that could be drawn from the information.

This distinction can seem obvious when examples are considered carefully, yet it is easily blurred in ordinary interaction. AI-generated responses can be fluent and confident. When a system responds in a way that appears to recognise something about the user, the distinction between what the user actually said and what the system derived from those statements may become less visible.

The system may appear to “know” something that was never directly disclosed.

A more precise description may be that information provided within the interaction supported an inference.

That distinction matters.

Inference Is Part of Everyday Reasoning

Inference is not unique to artificial intelligence.

Human beings constantly draw conclusions from incomplete information. Dark clouds may suggest rain. A closed shop during its normal opening hours may suggest an unexpected interruption. Repeated changes to a project schedule may suggest difficulties even when nobody has explicitly announced a problem.

Much of ordinary reasoning depends upon connecting observations and drawing conclusions that go beyond what has been directly stated.

AI systems can also operate through relationships and patterns within information, although the mechanisms involved differ fundamentally from human experience and vary among different forms of AI. The purpose here is not to explain those technical mechanisms. It is to recognise the informational consequence: AI can sometimes produce outputs that go beyond merely repeating the information supplied to it.

This ability contributes to usefulness. If an AI system could do nothing except reproduce the exact information it received, many forms of analysis, classification, interpretation and contextual assistance would be greatly limited.

Inference therefore should not automatically be understood as an undesirable feature.

At the same time, the ability to derive something new from existing information expands the conceptual scope of the information relationship.

Ordinary Information Can Support Less Obvious Conclusions

Inference also challenges the idea that information falls neatly into two categories: sensitive and non-sensitive.

An individual detail may appear ordinary. Occupation, approximate location, interests, routines or preferences may not seem particularly revealing when considered separately. But

combinations of information can sometimes support conclusions that none of those details independently establishes.

Consider someone discussing plans for the coming year. Over several exchanges they mention an expected change in working hours, questions about housing in another city, the timing of a school year and the availability of particular professional opportunities.

No statement says, “My family may relocate because of a change in my employment.”

Yet the combination could make such an interpretation plausible.

It could also be wrong.

The person may be researching for someone else. The questions may relate to several unrelated situations. Some information may be hypothetical. The user may simply be exploring possibilities.

This is one of the defining characteristics of inference: it can extend understanding while also introducing uncertainty.

The fact that information supports a conclusion does not transform that conclusion into an established fact.

Inference Can Concern Characteristics, Preferences and Intentions

What can potentially be inferred is not limited to factual circumstances.

Patterns of information may appear to indicate preferences. Repeated interest in particular subjects may suggest priorities. Language may appear to indicate familiarity with a topic. Sequences of questions may suggest intentions. Choices made during an interaction may contribute to an impression of what kinds of responses a person finds useful.

Again, “suggest” is doing important work.

A person who frequently asks for concise explanations may prefer concise responses, or they may simply be working under time pressure on those occasions. Someone repeatedly discussing a particular country may live there, plan to travel there, conduct business there, study it professionally or simply have an interest in it.

The observable information may be identical while the underlying explanation differs.

Inference therefore creates a separation between **what can reasonably be suggested by information** and **what is actually true about the individual**.

This Guide does not attempt to establish methods for verifying such conclusions. Truth and verification have their own distinct territory within the AIGuide series. The concern here is the information relationship itself: conclusions about a person may potentially emerge without that person ever having explicitly supplied those conclusions.

Inferences Can Generate New Information

This leads to a further conceptual question.

If an AI system derives a conclusion from information provided by a user, is the resulting conclusion itself information about the person?

In practical terms, it may function as such.

Suppose a person directly provides facts A, B and C. A system uses those facts to derive D. The person never supplied D, but D may now become part of the informational context within which subsequent interaction occurs.

This illustrates why an information relationship cannot always be understood simply by listing what the person originally entered.

The relationship may develop.

Information can lead to inference, and inference can potentially influence subsequent interpretation, personalisation or interaction. Depending on the system involved, a derived conclusion might be temporary and relevant only to the immediate context, or it might interact with broader forms of continuity or personalisation.

Those possibilities will be examined later. What matters here is the conceptual transition:

information does not merely describe; it can also provide the basis from which additional informational propositions are generated.

Accurate, Inaccurate and Uncertain Inferences

The possibility of inference introduces an important complication: derived information may be accurate, partly accurate, outdated, context-dependent or wrong.

A person may repeatedly ask about a health condition because they are supporting a relative rather than experiencing the condition themselves. Questions about senior management may come from a researcher rather than an executive. Requests concerning university admissions might come from a parent, teacher, counsellor or prospective student.

If the observable behaviour were treated as direct evidence of identity or circumstances, an inaccurate picture could result.

Even an inference that was once reasonable can cease to reflect current circumstances. Preferences change. Employment changes. Relationships change. Interests shift. A person's behaviour in one context may not represent their behaviour in another.

This creates a fundamental difference between a directly supplied fact and a derived proposition.

Directly supplied information can itself be inaccurate, of course. People can make mistakes, simplify circumstances, provide hypothetical scenarios or intentionally supply fictional material. But with inference there is an additional interpretative step between the underlying information and the resulting conclusion.

That step introduces another source of uncertainty.

The Human Tendency to Read Recognition as Knowledge

Conversational AI can make inference particularly noticeable because the resulting response is expressed in natural language.

If an AI response reflects a person's apparent preferences or circumstances, the experience can feel like recognition. The system appears to understand the individual.

That impression may be useful in creating a coherent interaction, but it can also blur several distinct concepts.

The system may be responding to information available in the immediate conversation. It may be drawing an inference from what the user has said. In some environments, broader contextual or personalised information may be relevant. Different systems can operate differently.

These possibilities should not be collapsed into the human concept of “knowing someone.”

Human knowledge of another person develops through social relationships, memory, interpretation and lived experience. Computational treatment of information is different.

The distinction becomes increasingly important as AI interactions become more fluent and personalised. A system that produces an apparently perceptive response can create the impression that it possesses a coherent understanding of the individual behind the conversation.

What appears as knowledge may sometimes be better understood as a combination of available information, contextual processing and inference.

The broader question of what it means to say that AI “knows” something about a person will be examined later in this Guide.

Inference Changes the Boundary of Disclosure

Inference also complicates the everyday idea of disclosure.

A person may reasonably say, “I never told the AI that.”

In a literal sense, that may be entirely correct.

But if other information was provided from which the unstated conclusion could be derived, the informational relationship may extend beyond the boundary the person associates with direct disclosure.

This does not mean the person should have anticipated every possible interpretation of everything they supplied. In complex information environments, doing so may be unrealistic.

It instead reveals why understanding data and AI requires more than identifying explicit fields, prompts or uploaded materials.

There are at least three layers now visible.

There is information deliberately provided.

There is information revealed through context, language and patterns of interaction.

And there is information potentially inferred from what is available.

These layers can overlap, but they are conceptually different.

An Information Relationship That Can Expand

Inference demonstrates why the information relationship between individuals and AI can become broader than the original exchange appears to suggest.

A user may begin with a small collection of facts. Those facts may contain contextual signals. Relationships among them may support interpretations. Those interpretations may create further informational possibilities.

None of this means that the resulting picture is complete or correct.

Indeed, inference creates a paradox at the centre of the information relationship: **an AI system may be capable of deriving information that a person never explicitly provided while simultaneously being wrong about what that derived information means.**

Both possibilities matter.

The first expands the potential informational reach of an interaction.

The second reminds us that informational reach is not the same as informational certainty.

Understanding inference therefore requires avoiding two opposite assumptions. One is that an AI system can know only what a person explicitly tells it. The other is that anything inferred from a person's information must accurately describe that person.

Neither assumption adequately captures the relationship.

Information can support conclusions. Conclusions can extend beyond direct disclosure. And those conclusions remain interpretations whose reliability may depend on the information, context, system and circumstances involved.

This brings the Guide to another dimension of the information relationship.

So far, the examination has focused largely on what exists within an interaction: what people provide, what they reveal and what may be inferred. But information also has a **temporal dimension**.

An interaction happens at a particular moment.

The information introduced into it does not necessarily have significance only at that moment.

The next question is therefore not simply what information exists, but what it means when information extends beyond the immediate interaction—when information, context or some aspect of the relationship **persists**.

Chapter 5 — When Information Persists

An AI interaction takes place at a particular moment. A person asks a question, uploads a document, provides context or continues a conversation. The system responds, and the immediate task may be completed.

From the user's perspective, that can feel like the natural end of the exchange.

But the end of an interaction and the end of the informational significance of that interaction are not necessarily the same thing.

The preceding chapters have examined what people deliberately provide to AI, what their interactions may reveal indirectly, and what may potentially be inferred from available information. These dimensions describe the informational content of an interaction. A further dimension concerns **time**.

What happens when information has relevance beyond the moment in which it was introduced?

This question leads to the concept of **persistence**.

Persistence does not describe one universal technical process. Different AI systems, services and operating environments can handle information in different ways. Information may exist for different periods, for different purposes and in different forms. Some interactions may be designed to be relatively self-contained, while others may support continuity across conversations or contribute to more personalised experiences.

The important distinction is therefore not that all information persists in the same way. It is that information should not automatically be assumed to lose all significance merely because the visible interaction has ended.

The Natural Mental Model of a Conversation

Human conversation provides an intuitive model for interacting with conversational AI.

A question is asked. An answer is given. The conversation moves forward or ends.

That familiar structure can influence how people understand the information involved. If an exchange feels temporary, the information within it may also feel temporary.

Yet even human conversations demonstrate that an interaction and its informational consequences are not identical. People remember what others have told them. A later conversation can be influenced by an earlier one. Information can alter how subsequent statements are interpreted.

AI systems are not people, and computational persistence should not be equated with human memory. But the conversational interface can encourage users to transfer human assumptions into a technically different environment.

A chat window may look ephemeral without establishing that every informational element associated with the interaction is ephemeral. Conversely, an interface that displays previous conversations does not by itself explain all of the ways information may be handled within or around the system.

The visible conversation and the wider information environment are therefore related but distinct.

Persistence Can Mean Different Things

The word *persistence* can itself be misleading if treated as though it describes a single state.

Information may persist in several conceptually different senses.

It may remain visible to the user as part of a conversation history. It may remain available as context during an ongoing exchange. Some systems may support features that allow certain information or preferences to influence later interactions. Information may also be handled for operational, safety, administrative or other purposes according to the particular environment in which a system operates.

These possibilities should not be assumed to apply identically to every AI service.

They also should not be collapsed into the simple statement that an AI “remembers everything.” Such language obscures more than it explains.

Persistence can be selective. It can be temporary. It can be purpose-dependent. What remains visible in an interface may differ from what is relevant technically or operationally. What influences one later interaction may not influence another.

Understanding persistence therefore begins by resisting an all-or-nothing model.

The useful question is not simply:

“Does the AI remember?”

It is:

“In what sense, for what purpose and for how long might information remain relevant?”

Continuity Changes the Experience of AI

Persistence can contribute to something users often value: continuity.

An AI system that can work with information already established within an extended interaction does not require every relevant detail to be repeated. The conversation can develop rather than continually restart. A system that supports some form of longer-term personalisation may similarly provide an experience that feels more adapted to the individual.

This illustrates why persistence should not automatically be framed as a problem.

Continuity can make interaction more efficient. It can reduce repetition. It can allow complex work to develop across multiple exchanges. It may enable AI assistance to respond more appropriately to established context.

At the same time, continuity depends upon information having significance beyond the instant in which it was first supplied.

The very feature that can make an interaction feel more coherent therefore also makes the temporal dimension of information more important.

A person may appreciate not having to explain the same context repeatedly while giving relatively little thought to what continuity requires informationally.

These are not contradictory positions. They are two aspects of the same relationship.

The Difference Between the User's Time Horizon and the Information's Time Horizon

When people provide information, they often think primarily about the immediate task.

A document is uploaded because it needs to be summarised now.

Personal circumstances are described because they help answer today's question.

Background information is provided because it improves the current conversation.

The user's **purpose horizon** may therefore be short.

The possible **information horizon** may be different.

This distinction matters because expectations about information are often shaped by purpose. If a person provides information for a task that lasts five minutes, they may intuitively associate the information with those five minutes.

But the duration of the purpose does not itself determine the duration or subsequent relevance of the information.

This does not mean that information necessarily continues indefinitely. It means that the lifespan of the task and the lifespan of the information should not be treated as conceptually identical.

A brief interaction can contain information whose significance extends beyond the interaction. A lengthy interaction may contain information that has little continuing relevance. Duration of use and duration of informational significance are separate questions.

Persistence Can Change Meaning

Time can also change what information means.

A statement such as “I work in Singapore” describes a circumstance at one point. Years later, it may no longer be accurate. A preference expressed during one project may not represent a person's general preference. A temporary family circumstance may later change. An employment role, educational status, location or personal objective can all become outdated.

This creates an important tension.

Persistence can provide continuity, but continuity can also carry forward information whose relevance has changed.

The problem is not unique to AI. Human records, organisational databases and personal recollections can all become outdated. AI interaction nevertheless makes the issue particularly interesting because information may be used contextually. If earlier information contributes to a later interaction, its continuing relevance can matter to how that interaction develops.

An informational relationship therefore has not only a history but also the possibility of **informational ageing**.

What was accurate may become outdated.

What was relevant may become irrelevant.

What reflected a temporary circumstance may appear more permanent when separated from its original context.

Persistence preserves information; it does not necessarily preserve its meaning.

Persistence and Inference Can Interact

The temporal dimension also connects with the inference examined in Chapter 4.

A conclusion drawn from one interaction may be tentative. If similar information appears repeatedly over time, the pattern may seem more significant. Conversely, an earlier inference may become less appropriate as circumstances change.

This illustrates how the different dimensions of the information relationship can interact.

Information can be provided.

Patterns can reveal additional context.

Inferences can be drawn.

Persistence can allow earlier information to remain relevant to later informational contexts.

None of these stages guarantees accuracy, permanence or consequence. But together they show why the information relationship can become more complex than a sequence of isolated prompts.

The relationship may acquire a history.

What the User Can See Is Not the Whole Concept

Users may understandably judge persistence by what they can see.

If a conversation appears in a history, persistence is visible. If an interface begins a new conversation with no obvious reference to previous exchanges, the experience may feel like a fresh start.

But interface visibility and informational persistence are not necessarily equivalent concepts.

Something visible may exist primarily for the user's convenience. Something not visible may still have relevance within a wider operational context. A system may distinguish among different forms of information, different purposes and different periods of availability.

This is one reason the information relationship can be difficult for ordinary users to understand.

The interface is necessarily a simplified representation of a more complex technical and organisational environment. A person can operate the interface successfully without possessing a complete mental model of how information moves through that environment.

The resulting gap between visible interaction and less visible information processes will receive fuller treatment later in this Guide.

For now, persistence provides a clear example of why that gap matters.

Persistence Is Not the Same as Permanence

It is equally important not to confuse persistence with permanence.

To say that information may extend beyond the immediate interaction is not to say that it exists forever.

Different information may have different lifecycles. Different systems may have different arrangements. Information can potentially be modified, separated from particular contexts, deleted, transformed or cease to be relevant, depending upon the circumstances involved.

The concept of persistence therefore describes **continuity through time**, not endless existence.

This distinction helps avoid two misleading extremes.

One is the assumption that closing a conversation necessarily ends every informational consequence associated with it.

The other is the assumption that anything ever provided to an AI system inevitably remains there permanently and unchanged.

Neither provides a sufficiently careful basis for understanding a diverse and changing AI environment.

From Persistence to an Accumulating Picture

Once information can extend across time, another possibility emerges.

Individual pieces of information may no longer stand alone.

A preference expressed today may sit alongside context provided previously. Repeated topics may create patterns. Earlier circumstances may influence how later requests are understood. Depending on the system and environment, continuity can contribute to an interaction that increasingly reflects information associated with the individual.

This is where persistence begins to connect with **personalisation**.

Personalisation can make AI interaction appear more relevant, efficient or individually responsive. But personalisation does not arise from nothing. To the extent that an experience reflects characteristics, preferences, history or context associated with an individual, information is part of its foundation.

The informational relationship can therefore develop from a momentary exchange into something with continuity.

That does not happen identically in every system, nor does it mean that every interaction builds a permanent profile of the person. The important point is that time changes what is conceptually possible.

Without persistence, each interaction can remain informationally isolated.

With some form of continuity, separate pieces of information may begin to contribute to a larger context.

The central question of this chapter is therefore not simply whether information is “kept.” It is what changes when information has relevance beyond the moment in which it was supplied.

Persistence can support continuity.

It can reduce repetition.

It can preserve context.

It can also carry forward information whose meaning has changed, connect interactions that the user may experience as separate, and make the informational history of an AI relationship increasingly significant.

The end of a conversation is therefore not necessarily the conceptual end of its information.

And once information can persist, the next question follows naturally:

What happens when that continuing information begins to shape the experience the AI provides?

That is the point at which persistence becomes part of a wider phenomenon—the development of an accumulating informational picture and the possibility of increasingly personalised AI interaction.

Chapter 6 — Personalisation and the Accumulating Picture

An AI system can respond to a person without knowing anything particular about that individual. A generic question can receive a generic answer, and for many purposes that may be entirely sufficient.

But AI interaction can also become more contextual.

A response may take account of information supplied earlier in a conversation. A system may adapt to preferences expressed by the user. In some environments, information from previous interactions or other available context may contribute to later exchanges. The resulting experience can appear increasingly tailored to the individual.

This is **personalisation**.

Personalisation is often discussed primarily in terms of usefulness. A system that has relevant context may avoid asking the same questions repeatedly. It may produce responses better suited to a person's circumstances, interests or preferred way of interacting. Greater continuity can make AI assistance feel less generic and more responsive.

Yet personalisation also has an information dimension.

For an interaction to become personalised, there must be some basis for distinguishing one context, preference, circumstance or person from another. That basis involves information.

The purpose of this chapter is not to assess whether personalisation is desirable or undesirable. Nor is it to compare how particular AI systems implement it. The more fundamental question is what personalisation reveals about the developing information relationship between individuals and AI.

Personalisation Depends on Difference

A completely generic system does not need to distinguish among individual users in order to provide the same general response to the same general request.

Personalisation introduces difference.

Two people may ask similar questions but receive differently framed responses because different contextual information is available. One person may prefer concise explanations; another may benefit from greater detail. One may be working within a particular professional environment; another may approach the same subject as a beginner. One person's previous conversation may establish context that another person's does not.

The more an AI interaction reflects these differences, the more information becomes relevant to the experience.

That information does not necessarily consist of a formal personal profile. It might simply be context available within the current conversation. It might include a preference explicitly stated moments earlier. In other settings, some form of continuity across interactions may contribute.

These arrangements can vary considerably among systems.

The conceptual point remains the same: **personalisation requires informational differentiation.**

Something must provide a basis for treating one interaction differently from another.

Personalisation Can Begin Within a Single Conversation

It is easy to associate personalisation only with information accumulated over long periods. Yet it can begin almost immediately.

Consider a person who asks an AI system to explain a complex concept. The first response is too technical, so the person says, “I am completely new to this subject. Please explain it without specialist terminology.”

The next response may be adapted accordingly.

The interaction has become more personalised even if no information survives beyond that conversation. The system is responding differently because it now has additional information about the user's stated level of familiarity and preferred form of explanation.

As the conversation continues, more context can accumulate. The person may identify what they already understand, which examples are useful and which aspects remain confusing. Each contribution can make later responses more closely aligned with the developing context.

This illustrates an important distinction.

Personalisation does not necessarily require a permanent record of a person. It can arise whenever available information allows an interaction to be adapted to the circumstances of that individual or that particular exchange.

Persistence, examined in Chapter 5, can extend this possibility across time. But context within a single interaction can already create an accumulating informational picture.

The Picture Is Built from Different Kinds of Information

The information contributing to personalisation may come from several sources.

Some may be explicit. A person may state that they prefer brief responses, identify their occupation, explain their level of knowledge or describe a particular objective.

Other information may emerge indirectly through interaction. Repeated requests for clarification may indicate that an explanation needs to be approached differently. A sequence

of questions can establish what has already been covered. Choices made during the conversation may reveal which direction the user wishes to pursue.

Inference can add another layer. Information available within an interaction may support tentative conclusions about what is likely to be useful or relevant, even when the person has not explicitly stated a preference.

These categories reflect the progression established in earlier chapters: information deliberately supplied, information indirectly revealed and information potentially inferred.

Personalisation can bring these informational layers together.

This does not mean every system uses all of them, nor that every personalised response depends upon an elaborate collection of personal information. In some circumstances, only a small amount of immediate context is necessary.

But as informational context becomes richer, the possibility of more differentiated interaction also increases.

An Accumulating Picture Is Not Necessarily a Complete Picture

Personalisation can create an impression that an AI system increasingly “understands” the person using it.

That impression requires care.

An accumulating informational picture is not the same thing as a complete understanding of an individual.

People are complex, contextual and changeable. A preference expressed in one setting may not apply in another. A person may want detailed analysis for professional work and very brief explanations for everyday questions. Someone may explore topics unrelated to their own circumstances. A user may deliberately test different viewpoints or construct hypothetical scenarios.

Information accumulated through AI interaction therefore represents only particular aspects of activity within particular contexts.

Even where a system can adapt effectively to previous information, the resulting picture may be partial.

It may also contain inaccuracies.

As Chapter 4 established, inference introduces uncertainty. As Chapter 5 showed, persistence can preserve information whose relevance changes over time. These two issues become particularly significant when accumulated information contributes to personalisation.

A system may have more information while still possessing an incomplete or outdated basis for understanding what is relevant now.

More information does not automatically mean more accurate understanding.

The Convenience of Not Repeating Ourselves

One of the clearest attractions of personalisation is the reduction of repetition.

Human beings generally do not expect to reintroduce themselves completely every time they continue an established relationship. Repeatedly explaining the same preferences, background and context can be inefficient.

AI interaction can create similar expectations.

If a person has already explained the purpose of a project, they may value being able to continue without restating it. If an interaction has established a preferred level of detail, retaining that context can make subsequent exchanges smoother. If an AI-enabled service can work with relevant continuing information, it may be able to provide assistance more quickly.

This convenience has an informational foundation.

The system can avoid asking again only if relevant information is already available in some form.

The connection between convenience and information is easy to overlook because users experience the benefit rather than the informational mechanism behind it. The visible result is less repetition. The underlying concept is continuity of context.

This does not make the convenience problematic. It simply shows that personalisation and information cannot be understood independently of one another.

When Accumulation Changes Significance

Individual pieces of information can also acquire different significance when they form part of a larger picture.

A person's preferred response length may seem relatively minor. So may their occupation, general interests, recurring projects or areas of concern. Each can be understood separately.

As more context accumulates, relationships among these pieces can become visible.

This does not mean that every collection of information becomes highly revealing. Nor does it mean that accumulation inevitably produces a detailed profile. But informational significance can change when separate elements are considered together.

The same principle applies beyond AI. A single fact often tells less than a connected set of facts.

Within AI interaction, however, the effect can become particularly noticeable because accumulated context may influence what happens next. Information is not merely present; it may contribute to how subsequent responses are framed or what appears relevant.

The accumulating picture can therefore become part of the interaction itself.

Later in this Guide, the broader significance of combining individually modest pieces of information will receive dedicated examination. Personalisation provides an early example of why accumulation matters.

Personalisation Can Shape What the User Experiences

As an interaction becomes more tailored, users may increasingly experience AI through an individually relevant informational frame.

This can affect examples, explanations, emphasis and the context used to interpret requests. The experience may consequently become less interchangeable with that of another user.

Again, this can be useful. Relevant context can improve coherence.

But personalisation also means that the user is no longer encountering only an abstract AI capability. They may be encountering that capability partly through information associated with their own interaction history or current circumstances.

This creates an important conceptual shift.

The question is no longer only:

“What can this AI system do?”

It can also become:

“What does this AI system do in the context of information associated with me?”

That is an informational question rather than a judgement about whether the resulting output should be trusted or how a person should make decisions from it. Those adjacent questions belong elsewhere in the AIGuide series.

Here, the concern is the informational foundation of the differentiated experience.

Personalisation Can Create Expectations

Once users experience continuity, they may begin to expect it.

A system that appears to recognise preferences in one context can create an expectation that it will recognise them elsewhere. If earlier information appears relevant in a later interaction, a person may assume other information is similarly available.

These assumptions may not always reflect how the particular system operates.

Different features may use different contexts. Different interaction modes may have different forms of continuity. Systems and settings may change. The information visible to the user may not provide a complete explanation of what is available in a particular interaction.

Personalisation can therefore produce not only an accumulating informational picture but also an accumulating set of **human expectations** about that picture.

Users may wonder what the AI retains, what it can draw upon and why it appears to recognise some things but not others.

The experience can become increasingly natural while the underlying information relationship becomes harder to describe precisely.

Personalisation and the Perception of Being Known

This leads to one of the most human dimensions of personalised AI.

When a system remembers context, adapts its explanations or appears to recognise recurring preferences, interaction can begin to feel less anonymous.

The person may experience a sense of familiarity.

That experience should not be dismissed. It is part of how human beings interpret responsive communication. But neither should familiarity be confused with a human relationship or human understanding.

The system's apparent familiarity has an informational basis.

What feels like “it knows how I like things explained” may reflect available information about previous preferences. What feels like “it understands my project” may reflect accumulated context. What appears to be recognition may involve information and inference rather than the kind of personal knowledge associated with human relationships.

This distinction becomes increasingly important because conversational fluency can make computational interaction feel socially familiar.

The richer the informational context becomes, the easier it may be to experience informational responsiveness as personal understanding.

From Personalisation to a Deeper Information Relationship

Personalisation demonstrates how the information relationship can evolve.

At first, a person may simply provide information to complete a task.

As interaction develops, contextual information can accumulate. Patterns may become visible. Inferences may become possible. Persistence can provide continuity. That continuity can support increasingly differentiated responses.

The information relationship is no longer merely about what was entered into a particular prompt.

It can become part of the character of the interaction itself.

This does not mean that every AI system develops a detailed picture of every user. It does not establish that personalisation is inherently beneficial or harmful. And it does not determine what information practices an individual should accept.

It establishes something more fundamental.

A personalised AI experience has an informational foundation.

The more an interaction appears adapted to an individual, the more important it becomes to understand what information contributes to that adaptation, how complete that information is, and how readily the user can distinguish informational responsiveness from human-like understanding.

That last distinction opens another dimension of the relationship.

As conversational AI becomes familiar, responsive and personalised, people may disclose information in an environment that feels remarkably like private conversation. The tone can be direct. The interaction may feel individual. The system may appear attentive to context.

But a conversation that **feels private** and a relationship that is **confidential** are not necessarily the same thing.

Understanding that difference is the next step in making the information relationship visible.

Chapter 7 — The Difference Between Conversation and Confidentiality

Conversation is one of the most familiar ways in which people exchange information.

When someone speaks with another person, the nature of the relationship often influences what they are prepared to say. A conversation with a close friend may contain information that would not be shared with a stranger. A discussion with a doctor, lawyer or other professional may take place within particular expectations and obligations. A workplace conversation may be shaped by organisational relationships and responsibilities.

Conversational AI introduces a different kind of interaction.

A person can type in ordinary language, ask follow-up questions, explain circumstances and receive responses that appear attentive to what has been said. As Chapter 6 examined, an AI interaction may also become increasingly contextual or personalised. The system can appear familiar with the subject under discussion and may respond in a way that feels individually relevant.

These characteristics can make the interaction feel conversational in a deeply familiar sense.

But the **experience of conversation** and the **existence of confidentiality** are not the same thing.

Understanding that distinction is an important part of understanding the information relationship between people and AI.

Why Conversation Can Feel Private

A conversational AI interface can create a strong sense of one-to-one interaction.

The user types a message. A response appears directly beneath it. There may be no other visible participant. The conversation can take place on a personal device, perhaps in a private room, at any time of day. The user may ask questions they have not asked another person or describe situations that would be difficult to raise elsewhere.

From the user's perspective, the experience can therefore feel private.

This perception is understandable. Human beings interpret communication partly through its form. A direct exchange between “me” and “you” ordinarily carries different social expectations from posting information publicly or addressing a large audience.

Conversational AI reproduces some of the surface characteristics of direct interpersonal communication without reproducing the human relationship that ordinarily gives those characteristics their social meaning.

The interface can therefore create an important distinction between **perceived privacy** and the wider informational circumstances of the interaction.

A conversation can feel private because nobody else appears to be present.

That feeling alone does not explain what confidentiality means in the environment through which the conversation occurs.

Confidentiality Is More Than the Absence of an Audience

In ordinary language, privacy and confidentiality are sometimes used interchangeably. Conceptually, however, they can describe different aspects of an information relationship.

A person may regard something as private because it concerns matters they do not ordinarily share widely.

Confidentiality concerns the conditions under which information is received, handled or expected to remain restricted.

These concepts can overlap, but they are not identical.

A diary entry can be highly private without involving a confidential relationship with another person. A professional may receive information under recognised confidentiality obligations even when the information itself would not ordinarily be considered especially personal.

AI interaction introduces further complexity because the apparent conversational partner is a computational system operating within a wider technical and organisational environment.

The person may experience a direct exchange with the AI. The information relationship, however, cannot necessarily be understood solely by looking at the visible participants in the conversation.

This Guide does not provide legal analysis of confidentiality, nor does it determine what legal or contractual duties apply to particular services, organisations or jurisdictions. Those questions depend upon circumstances outside its scope.

The narrower point is that **conversational form does not itself establish confidential status**.

The Language of Conversation Can Encourage Disclosure

The naturalness of conversational AI can influence the way people communicate.

Traditional digital forms often make information provision explicit. A user encounters separate fields for name, address, occupation or other details. The structure signals that information is being collected.

Conversation works differently.

Information can emerge gradually as part of explaining a problem.

A person may begin with a general question, receive a response, add context, clarify what happened and then provide further details as the exchange develops. No single moment necessarily feels like a significant act of disclosure.

Over several exchanges, however, a considerable amount of information may become part of the conversation.

This is not necessarily because the AI system asked for sensitive information or because the user consciously decided to create a detailed personal account. It can result from the ordinary dynamics of dialogue: explanation leads to clarification, which leads to further explanation.

The very quality that makes conversational AI accessible—its ability to communicate through ordinary language—can therefore make information provision feel less formal.

A conversation can accumulate information without ever resembling a form being completed.

Responsiveness Can Strengthen the Sense of Confidentiality

Personalisation can reinforce this effect.

When an AI response reflects details provided earlier, the system appears attentive. When it adapts to the person's preferred style or remembers relevant context within an interaction, the exchange can feel coherent. When responses acknowledge circumstances described by the user, the experience can resemble being listened to.

These characteristics can encourage a human interpretation of the interaction.

A person may begin to feel that the system “understands” the situation or “knows” what they mean.

As earlier chapters have emphasised, such language requires care. Informational responsiveness is not equivalent to human understanding. An AI system's ability to use available context does not transform the interaction into a human relationship.

The distinction is particularly important here because human expectations of confidentiality often arise from relationships.

Someone may trust a close friend because of the history between them. A professional relationship may carry established ethical or legal expectations. An organisational relationship may operate under defined policies and responsibilities.

An AI system's conversational responsiveness does not, by itself, create any of those foundations.

The interaction may feel relational while the underlying basis remains informational and computational.

The Wider Environment Behind the Conversation

The visible AI conversation is usually only one part of a wider environment.

There may be a service provider, software application, organisational account, technical infrastructure and various policies or processes surrounding the interaction. In workplace or institutional settings, additional arrangements may influence how the service is configured or used.

These circumstances can vary greatly.

That variation is one reason broad claims about “AI confidentiality” can be misleading. There is no single information environment shared by every AI interaction worldwide.

A person using a publicly available AI service may be operating in different circumstances from an employee using an AI capability within an organisation. A system operating locally may differ from one relying upon remote infrastructure. Different providers and jurisdictions may establish different arrangements.

The existence of these differences does not require the ordinary user to become a technical or legal specialist.

It does, however, reinforce a central theme of this Guide: **the interface does not necessarily reveal the whole information relationship.**

What looks like one person having one conversation with one AI system may sit within a considerably broader informational environment.

Confidentiality Can Also Concern Other People's Information

The distinction becomes particularly significant when a conversation includes information about someone other than the user.

A person may describe a colleague, family member, customer, patient, student or friend because that person's circumstances are relevant to the question being asked. They may paste correspondence written by somebody else or upload a document containing information about several people.

From the user's perspective, the purpose may remain straightforward: obtain help with a task.

But the conversational setting can make it easy to focus on the usefulness of the exchange rather than on the fact that information about another person has entered it.

This does not establish whether the information should or should not have been provided. Such conclusions can depend upon legal, professional, organisational and personal circumstances that are outside the prescriptive scope of this Guide.

It demonstrates instead that perceived confidentiality has implications beyond the person operating the AI.

If an interaction feels like a private conversation, the user may intuitively extend that feeling to information concerning other people as well.

The information relationship may therefore be wider than the apparent relationship between one user and one AI system.

A Private Setting Is Not the Same as a Private Information Environment

Another distinction concerns physical privacy.

A person may use AI while alone, on their own device, through an account only they access. Nobody nearby can see the screen or hear the interaction.

The physical setting is private.

That tells us something important about who can observe the user directly. It does not, by itself, describe how information is handled within the digital environment.

This distinction is familiar in other areas of online life. Sending a message while sitting alone does not mean that the communication exists only on the device in the person's hand. Using an online service privately does not mean that no technical or organisational processes exist beyond what appears on the screen.

AI interaction adds conversational realism to this established digital distinction.

The user can therefore experience two different forms of privacy simultaneously: **privacy of physical interaction** and a much more complex question concerning the **information environment**.

Confusing the first with the second can produce an incomplete understanding of the relationship.

Confidentiality Is Contextual

There is also no single expectation of confidentiality that can be applied to every type of information.

A casual question about a recipe is informationally different from a detailed description of a personal dispute. A publicly available report differs from an internal organisational document. A fictional scenario differs from information about identifiable people.

Similarly, the environments in which AI is used differ.

For this reason, confidentiality is better understood here as a contextual feature of information relationships rather than a simple label attached to AI as a whole.

The relevant context may include the nature of the information, the environment in which the AI operates, the people represented in the information, and the organisational or institutional setting surrounding its use.

This Guide does not convert those considerations into a decision rule.

Its purpose is to make visible why the question cannot be answered merely by asking whether the interaction *looks* private.

The Gap Between Feeling and Understanding

Conversational AI presents an unusual combination.

It can feel informal while involving potentially significant information.

It can feel private while operating within a wider technical environment.

It can feel attentive without possessing human understanding.

It can feel familiar without establishing the kind of human relationship from which expectations of confidentiality often arise.

None of these observations means that users are mistaken to experience AI conversationally. The conversational interface is real, and its accessibility is part of what makes contemporary AI useful.

The important distinction is that **human experience of the interface and informational characteristics of the system are different layers of the same interaction.**

Understanding one does not automatically provide complete understanding of the other.

This distinction becomes increasingly significant as people use AI for more personal, professional and context-rich purposes. The more natural the conversation becomes, the easier it may be for the informational mechanics of the interaction to fade into the background.

From Confidentiality to Consent

Once the distinction between conversation and confidentiality becomes visible, another question follows.

If the information relationship is broader and more complex than the conversational interface makes apparent, how fully can an ordinary person understand what they are participating in?

People routinely encounter notices, terms, choices and settings associated with digital services. Some information practices may be explained explicitly. Others may depend upon technical or organisational details that are difficult to translate into everyday understanding.

This creates a deeper question than whether someone clicked a button, accepted a term or voluntarily entered information.

It raises the question of what it means to **understand and consent to an information relationship** when that relationship may be complex, changing and only partly visible from the interface.

That question cannot be answered merely by observing that a person chose to use the system.

Nor can it be answered by assuming that users must understand every technical detail before meaningful interaction is possible.

Between those extremes lies one of the central challenges of the modern information environment.

Conversational AI makes that challenge especially visible because the interaction can feel so simple.

A person asks.

The AI answers.

The exchange may feel direct, private and familiar.

But the informational relationship surrounding that exchange can be more complicated than the conversation itself suggests.

Recognising the difference between **conversation and confidentiality** does not determine what people should disclose or which systems they should use.

It provides something that must come first: a clearer understanding that the feeling of a private conversation and the conditions governing information within that conversation are not necessarily the same thing.

The next question is therefore whether people can meaningfully understand the information relationship well enough to make choices about participating in it.

That takes the Guide from **confidentiality** to **consent**.

Chapter 8 — Consent in a Complex Information Environment

Consent is often expressed through a simple action.

A person selects “accept.” They enable a feature. They agree to terms. They choose to upload a document, continue a conversation or provide information because doing so allows an AI-enabled service to perform the task they want completed.

At the level of action, participation can appear straightforward.

At the level of understanding, the picture may be more complicated.

The preceding chapters have progressively widened the information relationship between individuals and AI. People deliberately provide information. Their interactions may reveal additional information. Further information may be inferred. Some information may persist beyond the immediate exchange. Accumulated context can contribute to personalisation. And a conversational interface that feels private does not, by itself, explain the conditions surrounding confidentiality.

Together, these dimensions create an important question: **what does it mean for a person to understand and consent to an information relationship that may be more complex than the interaction visible on the screen?**

This is not principally a legal question for the purposes of this Guide. Different jurisdictions establish different requirements concerning consent, privacy, data protection and information handling. Particular organisations and services may also operate under specific obligations. Those matters are outside the Guide's scope.

The concern here is more fundamental and human-centred: the relationship between **choice, information and understanding**.

Consent and Understanding Are Related but Different

A person can make a choice without understanding every process associated with that choice.

This is common in modern life. People use banking systems without understanding every stage of payment processing. They travel by air without knowing the engineering of the aircraft. They use mobile communications without understanding the complete telecommunications infrastructure that makes the connection possible.

Complete technical understanding is therefore not a realistic condition for ordinary participation in complex systems.

AI is no different in this respect.

A person should not need to understand machine-learning architecture, computational infrastructure or every information-processing mechanism simply to ask an AI system a question.

Yet there is an important distinction between not understanding **how a technology works technically** and not understanding **the nature of the information relationship in which one is participating**.

A person may not need to know how an AI system technically processes a document to recognise that the document contains information about them or others. They may not need to understand the engineering of personalisation to appreciate that continuing information can potentially shape later interactions.

The relevant question is therefore not whether the user possesses complete technical knowledge.

It is whether the user has a sufficiently meaningful picture of the information relationship to understand the character of the choice being made.

The Simplicity of the Interface Can Hide the Complexity of the Relationship

Modern digital services often make difficult processes appear simple.

That simplicity can be valuable. Good interfaces reduce unnecessary complexity and allow people to accomplish tasks without becoming specialists in the technology beneath them.

Conversational AI takes this accessibility particularly far.

There may be no complicated form. The person simply starts typing.

This ease can create an unusual contrast. The interaction becomes simpler while the information environment underlying it may remain complex.

A user can describe a personal situation in ordinary language, upload a substantial document or continue a contextual conversation without encountering anything that feels like a formal information-management process.

The interface says, in effect:

Talk naturally.

The informational significance of doing so may be much harder to express in equally natural terms.

This is not necessarily evidence that complexity is being deliberately concealed. Complex systems are difficult to represent fully through simple interfaces. Excessive explanation can itself become a barrier to comprehension.

The resulting tension is important: **making interaction easier does not necessarily make the information relationship easier to understand.**

Agreement Is Not the Same as Attention

Consent mechanisms also exist within environments in which people encounter large amounts of information.

Terms, notices, explanations, settings and choices compete for attention with the actual reason the person is using the service.

The user's immediate objective may be to obtain an answer, complete a task, analyse a document or solve a problem. Information notices are therefore encountered within a broader competition for attention.

A person may technically have access to an explanation without reading it closely.

They may read it without fully understanding it.

They may understand individual statements without being able to combine them into a coherent picture of the overall information relationship.

These are different states.

The distinction matters because the existence of information about information practices does not automatically establish that ordinary users have understood those practices.

Conversely, the fact that many users do not read every notice does not establish that information cannot be communicated meaningfully.

The challenge lies between those positions.

Consent occurs not in an abstract environment of unlimited time and attention but within the practical circumstances of human behaviour.

Complexity Can Be Distributed Across Many Decisions

Another difficulty is that the information relationship may not be created through one single decision.

A person might first create an account. Later they begin a conversation. In another interaction they upload a document. At some point they enable or encounter personalisation. Over time they provide additional context. New capabilities may become available. The service itself may change.

No single action necessarily represents the complete relationship.

Instead, the informational environment can develop through a series of choices.

This mirrors the accumulating picture examined in Chapter 6. Just as information can accumulate, so too can the decisions through which that information relationship develops.

A person may understand each action separately while having a less complete picture of their combined effect.

They know they uploaded one document.

They know they explained one personal situation.

They know they continued a previous interaction.

What may be less visible is how these separate moments relate to the wider informational context.

Consent can therefore have a temporal dimension just as information does.

Consent to One Purpose Does Not Explain Every Possible Meaning

People also make choices in relation to purposes.

A document is provided because the person wants it summarised. Context is supplied because it should improve an explanation. Personal information is introduced because it appears relevant to a question.

From the user's perspective, the purpose can define the informational transaction:

“I gave this information so that the AI could do this task.”

That is an understandable mental model.

But earlier chapters have shown why the informational meaning of an interaction can be broader than its immediate purpose. A document can contain information beyond the section being analysed. Interaction can reveal patterns. Information can support inference. Persistence can extend relevance through time.

The purpose for which information was provided and the wider informational characteristics of the interaction are therefore not necessarily identical.

Recognising this difference does not determine whether any particular additional use is appropriate. That question may involve legal, contractual, organisational or ethical considerations beyond this Guide.

It demonstrates why understanding consent requires attention not only to the act of providing information but also to the person's understanding of the relationship surrounding that act.

Consent Becomes More Difficult When Inference Is Involved

Inference creates a particularly interesting challenge.

A person can knowingly provide information A and information B.

But what does it mean to understand the possibility that A and B together may support inference C?

The person did not directly provide C.

They may not have anticipated C.

They may disagree with C.

C may even be wrong.

Yet the possibility of deriving additional information means that the informational significance of what was supplied cannot always be understood merely by listing the original inputs.

This makes exhaustive anticipation difficult.

People cannot reasonably be expected to imagine every conclusion that might potentially be drawn from every combination of information they provide. The possibilities can be too numerous, uncertain and context-dependent.

At the same time, ignoring inference entirely would leave an important part of the information relationship unexplained.

Meaningful understanding therefore sits between two unrealistic extremes: assuming that a person must foresee every possible informational consequence, and assuming that awareness of direct disclosure is sufficient to understand the entire relationship.

Consent Can Be Complicated by Information About Others

The issue becomes still more complex when the information concerns another person.

Someone may upload correspondence written by a colleague, describe a family member's circumstances or provide a document containing information about customers, employees or other individuals.

The person interacting with AI is making the immediate decision to provide the material.

The information represented within it may concern people who are not making that decision.

This separates two roles that can easily be treated as though they were identical:

the person providing the information and the person the information is about.

Sometimes they are the same person.

Sometimes they are not.

The legal, professional and ethical implications of such situations can vary widely and are not resolved here. The conceptual importance is that an information relationship can involve people who have never interacted with the AI system themselves.

Consent therefore cannot always be understood purely as a matter between one user and one AI service.

Consent Exists Within Unequal Understanding

Complex information environments can also produce differences in knowledge between participants.

The organisation operating a service may understand aspects of its technical and informational arrangements that an ordinary user does not. Specialists may understand distinctions that are difficult to communicate concisely. Users themselves differ greatly in digital familiarity, technical knowledge, language and experience.

This creates an informational asymmetry.

The existence of asymmetry does not automatically mean that users have no meaningful agency. Nor does it mean that organisations can eliminate every difference in understanding.

It means that the two sides of an information relationship may not begin with the same view of that relationship.

For the ordinary user, much depends upon what can be made visible and understandable without requiring specialist expertise.

This issue leads directly towards the visibility questions examined in the next chapter.

Consent Is Also a Continuing Question

AI systems and services do not remain static.

Capabilities develop. Features change. Personalisation can become more sophisticated. New forms of interaction can emerge. A person who initially used AI only for generic questions may later begin uploading documents, working with images or using increasingly contextual assistance.

The informational relationship can therefore change even if the person's original decision to begin using AI occurred much earlier.

This makes it useful to distinguish **initial participation** from **continuing understanding**.

An initial choice establishes that a person decided to enter an interaction under the circumstances existing at that time.

It does not necessarily establish that every later development is understood simply because the original choice was made.

Again, this observation does not prescribe how consent should be designed or renewed. Those questions can involve legal and policy considerations.

It identifies a more basic characteristic of the relationship: where information practices and AI capabilities evolve, understanding may also need to evolve if it is to remain aligned with the environment being experienced.

From Formal Choice to Meaningful Understanding

Consent in AI interaction is therefore more complicated than the presence or absence of an “accept” button.

Formal choices matter, but they do not describe the entire human experience of understanding an information relationship.

People operate with limited time and attention. Interfaces simplify complex systems. Information can accumulate gradually. Purposes can change. Inferences may extend beyond direct disclosure. Information can concern people other than the user. And the systems themselves can evolve.

None of this means meaningful consent is impossible.

It means that **choice and understanding should not be treated as interchangeable concepts**.

A person may choose to participate while possessing only a partial picture of the informational environment. The degree of understanding can vary from one individual to another and from one type of interaction to another.

This brings the Guide to a broader problem.

Much of the information relationship examined so far is only partly visible from the user's immediate perspective.

People can see what they type.

They can usually see what they upload.

They can see the AI's response.

But inference, persistence, accumulation, personalisation and the wider information environment may be considerably harder to observe directly.

The user therefore encounters AI through a visible interface while participating in an information relationship whose full dimensions may extend beyond what that interface reveals.

That difference between **what can readily be seen** and **what may be occurring within the broader information relationship** is the next subject of this Guide.

It is the **visibility gap**.

Chapter 9 — The Visibility Gap

When a person interacts with an AI system, some parts of the information relationship are immediately visible.

The user can see the words they type. They know when they upload a document or submit an image. They can read the AI's response and observe how the conversation develops. If previous exchanges remain displayed, they can see that history. If the system refers to information provided earlier, they may recognise that some form of continuity is influencing the interaction.

Other dimensions are much harder to observe.

A person may not be able to see what can be inferred from the information they provide. They may not know intuitively which information remains relevant beyond the immediate interaction. They may experience personalisation without having a complete picture of the information supporting it. They may understand that an AI service operates within a wider technical environment while having little direct visibility into that environment.

This difference between the interaction people can readily observe and the broader information relationship they may be participating in can be described as the **visibility gap**.

The visibility gap does not mean that everything beyond the interface is hidden, nor does it imply deliberate concealment. Information may be available through explanations, policies, settings or other mechanisms. Rather, the concept describes a more fundamental difficulty: **the information relationship can be more complex than what can be directly perceived through ordinary use of the system.**

The Interface Shows the Interaction, Not Necessarily the Whole Relationship

Interfaces are designed to make systems usable.

A conversational AI interface may present a text box, an upload function, previous exchanges and a small number of controls. This simplicity allows people to interact with sophisticated technology without understanding the technical processes operating beneath it.

That is one of the strengths of an effective interface.

But an interface is also necessarily selective. It shows what people need in order to interact with the system. It does not usually attempt to represent every technical, organisational or informational process associated with that interaction.

This creates an important distinction.

The interface can provide a clear representation of **what the person is doing** while providing a much less complete representation of **the wider informational environment in which that action occurs.**

A user can therefore become highly competent at operating an AI system without necessarily developing an equally detailed understanding of its information relationship.

Operational familiarity and informational visibility are not the same thing.

What We Provide Is Usually the Most Visible Layer

The most visible information is generally what the user consciously introduces.

If someone types their occupation into a prompt, they know they have provided it. If they upload a report, they can see the document being attached. If they submit an image, the act of providing it is explicit.

Even here, however, visibility has limits.

As Chapter 2 examined, a person may focus on the purpose of a document without noticing every piece of information it contains. Someone submitting a screenshot may concentrate on the portion relevant to the question while overlooking details elsewhere in the image.

The information is visible in principle, but not necessarily salient to the user at the moment of interaction.

Visibility is therefore not simply a binary condition in which information is either visible or invisible.

Something can be physically present on a screen yet receive little attention.

Something can be consciously supplied while its wider informational significance remains unclear.

This distinction becomes even more important as the Guide moves beyond direct provision.

Indirect Revelation Is Harder to See

Chapter 3 examined information that can be revealed through language, context and patterns of interaction.

These forms of information are less visible because they do not necessarily appear as separate items.

A user can see the individual questions they ask. They may be less likely to think of the sequence of those questions as an informational pattern.

They can see the words used in a prompt. They may not consider what those words suggest about familiarity with a subject, interests or circumstances.

They can recognise that they have provided context. They may not recognise every relationship among the details contained within that context.

The interaction therefore contains a distinction between **seeing the components** and **seeing what the components collectively reveal**.

This is one source of the visibility gap.

The user sees the conversation from the perspective of purpose: “These are the questions I needed answered.”

An informational perspective may also see a pattern: “These questions collectively describe an area of sustained attention.”

Both descriptions can be accurate. The second is simply less obvious from the user's immediate experience.

Inference Widens the Gap Further

Inference introduces an even greater separation between visible input and possible informational meaning.

A person can see what they told an AI system. They cannot necessarily see every conclusion that might be drawn from it.

Suppose a user provides several facts about their work, location and future plans. Each fact is visible because the user deliberately supplied it. But a conclusion derived from their combination may never have been stated.

The user sees the ingredients.

They may not see every interpretation those ingredients could support.

As Chapter 4 established, such interpretations are not necessarily correct. Inferences can be uncertain, incomplete or wrong. The issue here is not their accuracy but their visibility.

An inference may exist conceptually beyond the user's direct disclosure.

This means that the informational picture associated with an interaction can potentially become wider than the picture the person consciously remembers creating.

The visibility gap therefore grows as the relationship moves from direct information towards derived information.

Time Makes Visibility More Difficult

Persistence adds another challenge.

People experience interactions sequentially. A conversation occurs today. Another occurs next week. A preference may be expressed in one context and later appear relevant in another.

The longer the informational relationship extends through time, the harder it may become for a person to maintain a complete mental picture of everything that has contributed to it.

This is partly a human limitation.

People do not remember every detail they have entered into every digital service. They may forget what was said in an earlier AI conversation, what document was uploaded months ago or what preference was expressed in a particular context.

The information relationship may therefore become less visible not because the information was never observable, but because it has accumulated across time.

Persistence creates history.

History creates a further challenge: understanding the present may require remembering the past.

If an AI interaction reflects earlier context, the user may recognise the continuity. But they may not always know which particular information is relevant, how current it is or why one piece of context appears influential while another does not.

This is where persistence and personalisation begin to reinforce the visibility gap.

Personalisation Can Be Visible While Its Foundations Are Not

Personalisation creates an interesting contrast.

Its effects may be highly visible.

The AI appears to use a preferred style. It refers to an established project. It adapts an explanation to circumstances previously discussed. The person notices that the experience is no longer entirely generic.

But the informational basis of that personalisation may be less obvious.

Which information contributed?

Was it stated directly?

Was it inferred from the conversation?

Is it relevant only within the present interaction, or does some form of continuity extend further?

Is the information still accurate?

These questions may not have obvious answers from the personalised response itself.

The user can therefore observe the **effect** of informational differentiation without necessarily seeing its complete **informational foundation**.

This distinction matters because personalised interaction can feel intuitive. If the response fits the person well, there may be little reason in the moment to ask how that fit was achieved.

Usefulness can reduce the salience of the processes that make usefulness possible.

The Wider Environment Is Difficult to Represent

The visibility gap also exists because an AI service is not simply the conversational exchange displayed on a screen.

Depending on the system and context, there may be technical infrastructure, service providers, organisational arrangements, account settings, information-handling processes and other components surrounding the visible interaction.

For an ordinary user, constructing a complete model of this environment can be difficult.

Technical explanations may require specialist knowledge. Legal or policy documents may address issues using terminology unfamiliar to some readers. Organisational arrangements may differ across products, account types or jurisdictions. Features can change over time.

None of this means that information about the environment is necessarily unavailable.

It means that **availability and practical visibility are different concepts**.

Information can exist somewhere within documentation while remaining difficult for an ordinary person to integrate into a coherent understanding of the relationship.

This distinction connects directly with the consent issues examined in Chapter 8. A person can be provided with information and still face difficulty turning separate explanations into a usable mental model of what their participation means.

More Transparency Does Not Automatically Mean More Understanding

The obvious response to a visibility gap may appear to be simply providing more information.

Sometimes additional information can clearly improve understanding.

But quantity alone does not resolve the problem.

A person presented with extensive technical descriptions, multiple notices and detailed policies may technically have access to more information while becoming less able to identify what matters to their particular interaction.

Transparency and comprehension are therefore related but not identical.

Too little information can leave important features unexplained.

Too much poorly structured information can overwhelm.

Highly simplified information can omit significant distinctions.

Highly technical information can be inaccessible to non-specialists.

This is not a problem unique to AI. It is a recurring challenge in complex digital environments. AI makes it particularly significant because the visible interaction can be exceptionally simple while the informational relationships surrounding it can be comparatively sophisticated.

The challenge is therefore not merely to make information available, but to make complex relationships intelligible.

This Guide does not prescribe how providers should achieve that. Its purpose is to identify why the challenge exists.

Visibility Also Depends on the Observer

There is no single visibility gap experienced identically by everyone.

A technically experienced user may understand aspects of an AI system that another user does not. Someone familiar with digital information practices may interpret a feature differently from someone encountering AI for the first time. A professional operating within an organisation may have access to internal explanations unavailable to a general consumer.

Language, education, experience, accessibility and familiarity with technology can all affect what a person understands from the same interface or explanation.

Visibility therefore has a human dimension.

The question is not simply whether something has been disclosed somewhere.

It is also whether the person can reasonably perceive and understand its significance within the interaction they are having.

This reinforces the human-centred character of the information relationship. Technical facts matter, but the relationship is experienced through human attention, comprehension and expectations.

The Gap Can Never Be Eliminated Completely

Complex systems inevitably contain processes that ordinary users do not directly observe.

A person does not need to see every network connection, computational operation or infrastructure component in order to use AI meaningfully. Complete visibility would neither be practical nor necessarily useful.

The objective of understanding the visibility gap is therefore not to suggest that every underlying process must become continuously visible.

It is to recognise that **what people see should not automatically be treated as the complete informational reality of the interaction.**

That recognition helps avoid two opposite assumptions.

The first is that anything not visible must be problematic.

The second is that anything not visible can safely be ignored.

Neither follows.

Less visible processes may be routine, necessary and appropriately managed. They may also raise questions that matter to an individual's understanding of the information relationship. Their significance depends upon context.

Seeing the Relationship More Clearly

The visibility gap brings together much of what the Guide has examined so far.

People see what they type, but may overlook what their words reveal.

They see individual interactions, but may not perceive patterns across them.

They know what facts they supplied, but cannot necessarily see what might be inferred.

They experience conversations in the present, while information may have a longer temporal dimension.

They notice personalisation, but may not know exactly what information supports it.

They experience a conversational interface, but the wider information environment extends beyond the conversation itself.

None of these observations requires assuming that the unseen dimensions are harmful. The point is simpler: **the information relationship is larger than the portion of it that is immediately visible to the user.**

Recognising that gap provides a more realistic basis for understanding AI interaction.

But it also raises another question.

So far, much of the Guide has examined different kinds of information separately: a prompt, a preference, a behavioural pattern, an inference, a piece of continuing context.

Information rarely exists only as isolated pieces.

Several individually ordinary details can sometimes become more revealing when connected. Context can transform significance. A collection of small informational fragments can create a picture that none of them could create alone.

The next stage of the Guide therefore moves from visibility to **aggregation**.

The question is no longer simply what each piece of information says.

It is what may happen **when small pieces become a larger picture**.

Chapter 10 — When Small Pieces Become a Larger Picture

A single piece of information can appear unremarkable.

A person mentions the city in which they live. In another conversation, they identify their occupation. Elsewhere, they describe a hobby, a regular commitment, a preferred way of working or a particular concern. Each detail, considered separately, may seem to reveal relatively little.

Together, the picture can be different.

Information does not always derive its significance from individual pieces considered in isolation. Its meaning can change when separate details are connected, compared or placed within a wider context. Information that appears ordinary on its own may contribute to a considerably richer understanding when combined with other information.

This process can be described as **aggregation**.

Aggregation is not unique to artificial intelligence. Human beings routinely assemble fragments of information into larger pictures. Organisations combine records to understand activities over time. Researchers examine multiple observations to identify patterns. Everyday reasoning often depends upon connecting facts that individually provide only part of an explanation.

Within AI interaction, however, aggregation deserves particular attention because the information relationship can develop progressively. As earlier chapters have shown, people may deliberately provide information, reveal information indirectly, generate patterns through repeated interaction and create the basis for inference. If information also persists or contributes to personalisation, separate informational fragments may become relevant to one another.

The resulting picture may communicate more than any single piece from which it was formed.

Information Has Context

The significance of information depends partly upon context.

Knowing that someone regularly travels tells us one thing. Knowing where they travel, how frequently, for what apparent purposes and alongside what other activities can create a different picture.

Knowing a person's occupation provides one category of information. Combining occupation with location, seniority, professional interests and descriptions of current projects may provide considerably more context.

The individual pieces have not necessarily changed.

What has changed is their relationship to one another.

This illustrates an important characteristic of information: **meaning can emerge from combination**.

The effect can work in the opposite direction as well. Information that appears highly significant when isolated may become less revealing when additional context shows that an initial interpretation was incomplete or mistaken.

Aggregation therefore does not simply mean that more information always produces a more revealing or more accurate picture. It means that information can change meaning when considered alongside other information.

Context can strengthen an interpretation.

It can weaken one.

Or it can create an entirely different interpretation.

Small Details Can Become Informationally Significant

People often make intuitive distinctions between information they regard as important and information they regard as trivial.

Formal identifiers may immediately appear significant. Everyday preferences may not.

Yet significance can be cumulative.

Imagine that a person never provides a full personal profile in any single AI interaction. Instead, over time they mention their general location, type of work, family circumstances, preferred activities, recurring projects and plans for the future.

No individual interaction necessarily feels highly revealing.

Collectively, the information may describe a much more detailed set of circumstances.

This does not mean that every collection of ordinary details identifies a particular individual or produces a meaningful profile. Nor does it establish that any particular AI system actually combines information in this way.

The conceptual point is that the informational significance of a detail cannot always be assessed without considering the wider context in which that detail exists.

A piece of information that appears modest when viewed alone may become more consequential when connected to other pieces.

Aggregation Can Occur Within One Interaction

The larger picture does not necessarily require months or years of accumulated information.

It can emerge within a single conversation.

A person might begin by describing a workplace problem. As the AI asks for, or the person voluntarily provides, additional context, the conversation may gradually include the person's role, reporting relationships, organisational circumstances, recent events and intended objectives.

No single response creates the entire picture.

The picture develops through the interaction.

This is one reason conversational AI can create informational richness so quickly. Natural dialogue encourages contextual elaboration. A person can move easily from a general description to increasingly specific details because each addition appears relevant to obtaining a better response.

By the end of the exchange, the informational context may be considerably richer than it appeared at the beginning.

Aggregation therefore need not involve a deliberate act of assembling a profile. It can arise simply because multiple pieces of relevant context become part of the same interaction.

Aggregation Can Also Occur Across Time

Persistence introduces a further possibility.

Where information from separate interactions has continuing relevance, aggregation can extend across time.

A person's circumstances may emerge gradually rather than through one comprehensive disclosure. One interaction concerns work. Another concerns travel. Another concerns education or family circumstances. Another establishes preferences about how assistance should be provided.

If some of those informational elements become relevant together, a broader picture may emerge.

Again, systems differ significantly in what information can carry across interactions and how such continuity operates. The purpose here is not to assume a universal architecture.

It is to recognise the relationship between two concepts already established in this Guide.

Persistence allows information to extend through time. Aggregation allows separate information to acquire meaning through combination.

Where both occur, the information relationship can become cumulative.

Aggregation and Inference Reinforce One Another

Aggregation also interacts closely with inference.

A single piece of information may support only a weak conclusion. Several related pieces may make a particular inference appear more plausible.

Suppose a person asks one question about a foreign city. Very little can reasonably be concluded about why.

If subsequent interactions concern employment there, housing, transportation, schools and administrative arrangements, a possible interpretation begins to emerge: perhaps the person is considering relocation.

But even then, the conclusion remains an inference.

The person might be researching for someone else. They might be preparing a report, advising a client or developing a fictional scenario.

More information can strengthen an inference without turning it into certainty.

This is an important limit on the apparent power of aggregation.

A larger picture can still be wrong.

Indeed, aggregation can sometimes create misplaced confidence. When multiple details appear to fit a coherent explanation, the explanation can seem convincing even though an alternative interpretation remains possible.

The distinction established in Chapter 4 therefore remains essential: information can support conclusions without proving them.

A Larger Picture Can Contain Errors

Aggregation can also magnify problems in the underlying information.

Suppose several pieces of information are outdated. When considered separately, each may have limited effect. When combined into a larger picture, they may create an apparently coherent description of a person that no longer reflects their circumstances.

Similarly, an incorrect inference can interact with accurate information. If the inference is then treated as part of the wider informational context, later interpretations may develop around a mistaken assumption.

This means that an accumulating picture is not automatically a progressively accurate picture.

Information can accumulate faster than understanding.

A system may have access to more context while some of that context is incomplete, outdated, hypothetical or misunderstood.

Human beings experience the same difficulty. Knowing many things about someone does not guarantee understanding that person accurately. The meaning of information depends upon relevance, context and interpretation as well as quantity.

AI does not remove this distinction.

Aggregation Can Blur the Idea of Sensitive Information

The larger-picture effect also complicates simple classifications of information.

Some categories of information are readily recognised as personal or sensitive in everyday life. Other details appear routine.

Aggregation demonstrates why that distinction may not always capture the full informational significance of an interaction.

Several ordinary facts can sometimes reveal a circumstance that none states explicitly. A combination of interests, routines, location and contextual details may support conclusions about matters that the person never directly described.

The issue is not that every ordinary fact should therefore be treated as extraordinary.

That would make the concept too broad to be useful.

The more careful observation is that **informational significance is relational**. It depends not only on what a particular fact says but also on what other information surrounds it and what conclusions the combination may support.

This is one reason the question “Did I provide anything important?” can be harder to answer than it initially appears.

Importance may not reside in one item.

It may emerge from the whole.

The User May Never See the Whole Picture

Aggregation connects directly with the visibility gap examined in Chapter 9.

People experience information provision sequentially.

They remember today's question more readily than one asked months earlier. They may know what they have said within a current conversation but have only a general recollection of information supplied across many previous interactions.

As a result, the larger informational picture may be less visible to the individual than the pieces from which it was formed.

This is a distinctive human limitation.

The person lives the experiences that generate the information, but does not necessarily maintain a continuously updated inventory of every detail they have communicated.

The user may therefore know every individual fact and still not perceive the combined informational picture.

Visibility and aggregation are closely connected for precisely this reason.

The more dispersed the information is across contexts and time, the harder it may become for a person to recognise what the collection communicates as a whole.

Aggregation Does Not Require a Human-Like Profile

The idea of a “larger picture” can easily encourage anthropomorphic interpretations.

It may sound as though an AI system develops a complete portrait of a person in the way a close acquaintance gradually comes to know them.

That is not the claim being made.

Information can be aggregated computationally without producing anything resembling human understanding. A system can work with relationships among pieces of information while lacking lived experience, social knowledge of the individual or the broader understanding associated with a human relationship.

Nor should aggregation automatically be described as the construction of a permanent profile. Different systems can use information in different ways, for different periods and within different contexts.

The concept is narrower.

When separate pieces of information become available together, their combined informational meaning can differ from their meaning in isolation.

That principle holds whether the combination exists briefly within one conversation or forms part of a more persistent context.

The Larger Picture Can Include Other People

Aggregation becomes still more complex when the information concerns more than one individual.

A person's AI interaction may contain information about colleagues, family members, customers, friends or other people. A document may describe relationships among several

individuals. A conversation may gradually establish who did what, who is connected to whom and how different circumstances relate.

The resulting informational picture may therefore extend beyond the user.

This possibility has already appeared at several points in the Guide, but aggregation makes its significance clearer.

A single reference to another person may reveal little. Repeated references, documents, contextual explanations and descriptions of relationships can collectively communicate considerably more.

The information relationship can consequently become relational in another sense: it may contain information about **relationships between people**, not merely separate facts about individuals.

That territory requires its own examination.

From Pieces of Information to People Around Us

The central lesson of aggregation is simple but significant.

Information cannot always be understood one piece at a time.

A prompt may appear modest.

A preference may appear ordinary.

A contextual detail may appear incidental.

An individual interaction may appear temporary.

Yet when information is connected, its significance can change.

This does not mean that aggregation is inherently harmful. Combining information can be essential to producing coherent, contextual and useful AI assistance. Without connecting information, systems would often struggle to understand even the immediate subject of a conversation.

Nor does aggregation guarantee accuracy. A larger picture can remain incomplete, outdated or mistaken.

The point is that **the whole can communicate something different from the individual parts**.

This adds another layer to the information relationship developed throughout this Guide. People do not merely provide isolated pieces of information. Those pieces can exist within contexts, relate to one another, persist across time and support wider interpretations.

And the information being assembled does not necessarily concern only the person interacting with AI.

Often, people's lives are described through relationships with others. Their documents contain other people's names. Their questions concern colleagues, relatives, customers or friends. Their photographs include other people. Their stories describe events involving several participants.

Once small pieces become a larger picture, an important question therefore moves to the foreground:

Whose information is actually contained within that picture?

The answer may extend well beyond the person using the AI.

That is the subject of the next chapter: **Other People's Information.**

Chapter 11 — Other People's Information

When a person interacts with AI, it is natural to think of the information relationship as existing between that individual and the system.

The person provides information. The AI processes the interaction. The resulting questions about disclosure, inference, persistence and personalisation therefore appear to concern the user.

But people's information rarely exists in complete isolation.

A workplace document may contain the names and activities of colleagues. An email includes words written by somebody else. A photograph can show friends or family members. A person describing a disagreement necessarily provides information about the other people involved. Meeting notes may identify participants and record their contributions. A request for help understanding a situation can contain detailed information about someone who has never interacted with the AI system at all.

The information relationship can therefore extend beyond the person operating the AI.

This creates an important distinction:

the person providing information is not necessarily the person the information is about.

That distinction has appeared at several points earlier in this Guide. It now requires closer examination because it changes the apparent boundaries of “your data and AI.” What a person gives AI can include not only their own information but information belonging to, describing or involving other people.

Our Information Is Often Interconnected

Human lives are relational.

People work with others, communicate with others, live with others and make decisions that involve others. Their documents, photographs, correspondence and memories therefore contain information about relationships as well as individuals.

Consider an ordinary email conversation.

A person who copies an email into an AI system for help preparing a response may be providing their own words, but also the sender's words. The message may identify other people, describe previous events or contain information about an organisation.

The immediate task may be simple: “Help me reply to this.”

Informationally, however, the interaction can involve several people.

The same principle applies to documents. Meeting minutes may contain the views of multiple participants. A project report can describe employees, customers or suppliers. A school document may contain information about students and teachers. A family document can concern several relatives.

Information about one person is frequently embedded within information held by another.

AI does not create this interconnectedness. It makes the interconnectedness significant in a new context because material created for one purpose can now be introduced into an AI interaction for another.

People Naturally Describe Other People

Documents are not required for this to happen.

Ordinary conversation constantly contains information about others.

Someone may ask:

“My colleague keeps changing the project requirements. Can you help me understand how to structure the issues before our next meeting?”

The user's purpose concerns their own preparation. Yet the description also contains information about the colleague.

Another person may explain:

“My daughter is considering several university courses and is uncertain which subjects interest her most.”

Again, the person interacting with AI has introduced information about somebody else.

A customer may describe an employee. An employee may describe a manager. A friend may describe another friend. A professional may discuss circumstances involving clients or colleagues.

These interactions can range from vague descriptions to highly detailed accounts.

The important point is not that mentioning another person automatically creates a significant information issue. Human communication would be impossible if people could never refer to others.

The point is that **information does not become the user's own information merely because the user is the person who entered it into the AI system.**

The informational subject and the informational provider can be different people.

Other People May Not Know the Interaction Exists

This distinction creates an unusual feature of the information relationship.

The user knows they are interacting with AI.

The other person may not.

A colleague whose email is included in a prompt may have no awareness that the email has been introduced into an AI interaction. A person visible in a photograph may not know that the image is being analysed. Someone discussed within a personal account may never know that their circumstances formed part of a conversation with AI.

This observation does not determine whether such an interaction is appropriate. Circumstances vary too widely for a general conclusion. Legal, professional, contractual and organisational obligations may also apply in particular settings, and this Guide does not provide advice on those matters.

What the observation demonstrates is that **participation and representation are different**.

A person can become represented within an AI information relationship without personally participating in the interaction.

This expands the conceptual boundary of the relationship beyond the visible user.

Documents Can Carry Hidden Social Context

Uploaded materials make this particularly easy to overlook.

When someone sees a document primarily as material to be summarised or analysed, the people represented within it can fade into the background.

Consider meeting minutes.

To the person uploading them, they may be “the minutes.” But informationally, they may contain names, roles, statements, decisions, disagreements and responsibilities associated with several individuals.

Similarly, a chain of emails may be viewed as “correspondence,” while containing a detailed record of who communicated what, when and in response to whom.

A photograph may be “a picture from an event,” while depicting several identifiable people and aspects of their surroundings.

The task-oriented view and the information-oriented view therefore differ.

The task asks:

“What do I want the AI to do with this material?”

The information perspective also asks:

“Who and what is represented within this material?”

Both questions describe the same interaction from different directions.

One Person's Information Can Reveal Another's

Interconnected information can go further.

A person may provide information about themselves that simultaneously communicates something about another person.

If someone states that they care full-time for an elderly parent, the statement concerns the user but also reveals something about the parent's circumstances.

If an employee describes their reporting relationship and a particular workplace dispute, the account can communicate information about their manager or colleagues.

If someone discusses inherited characteristics within a family, information about one family member may have implications for others.

This illustrates why personal information cannot always be separated neatly by individual.

Human relationships create informational connections.

The same fact may concern several people simultaneously.

This does not mean that information about one person automatically establishes facts about another. As earlier chapters have shown, inference introduces uncertainty. But relational information can provide a basis for conclusions extending beyond the individual who originally supplied it.

The larger informational picture can therefore include people who never directly contributed a single word.

Aggregation Can Deepen Third-Party Information

Chapter 10 examined how individually modest pieces of information can become more revealing when combined.

The same principle applies to information about other people.

A single reference to a colleague may communicate very little. Repeated descriptions can gradually establish that person's role, responsibilities, working relationships and involvement in particular events.

A family member may be mentioned incidentally in one conversation and become part of a much richer contextual picture across several discussions.

Documents can accelerate this process because one upload may contain many connected details about multiple people.

Aggregation therefore changes the significance of third-party information just as it changes the significance of information about the user.

The resulting picture may still be incomplete or inaccurate. The user is describing events from their own perspective. Documents may be outdated. Interpretations may be contested.

More information does not eliminate these limitations.

It does, however, increase the amount of informational context in which other people can be represented.

Inference Can Extend Beyond the User

Inference adds another layer.

Suppose someone repeatedly describes difficulties involving a colleague. The explicit information concerns particular events. A broader conclusion about the colleague's behaviour or intentions would be an inference.

The colleague has not supplied the information.

They have not explained their perspective.

They may not even know the conversation exists.

Yet conclusions concerning them may emerge from information supplied by somebody else.

This demonstrates an important asymmetry.

Inference about a user can at least arise within an interaction in which that user is participating. Inference about another person may arise without that person's involvement at all.

The conclusion may be accurate, partly accurate or completely mistaken.

This is not principally a verification question for this Guide. The important point is informational: **derived information can concern people other than the person from whom the underlying information originated.**

The informational reach of an AI interaction can therefore extend beyond direct participants.

Professional and Organisational Contexts Add Complexity

The distinction can become especially important when AI is used in professional settings.

Work is information-rich.

Employees encounter internal documents, customer information, correspondence, meeting records, financial material, project information and other forms of organisational knowledge. Professionals in different fields may work with information concerning clients, patients, students, citizens or members of the public.

Introducing such material into AI interaction can therefore involve information about many people.

The significance of this varies greatly according to context. Different professions, organisations and jurisdictions operate under different requirements. Particular AI environments may also be designed and governed differently.

For these reasons, AIGuide 12 does not attempt to establish universal rules for professional use.

Its role is to identify the conceptual issue beneath those differences: when AI is used to work with information, the information relationship may involve **multiple human subjects even though only one person is operating the system**.

That observation is relevant across national and professional boundaries.

Family Information Requires a Clear Boundary

Family relationships provide another obvious example because information about relatives is often deeply interconnected.

A person may discuss a partner, parent, child or sibling when seeking help understanding a situation. Family documents and photographs can contain information about several people. One person's circumstances may reveal something about another's.

These examples belong within AIGuide 12 only to the extent that they illustrate **third-party information**.

The broader questions surrounding AI within family relationships—such as roles, responsibilities, household boundaries and the particular implications of AI in family life—belong primarily to AIGuide 11, *AI in the Family*.

Maintaining that distinction is important.

Here, a family member is an example of a wider informational principle that applies equally to colleagues, customers, friends, students and others: **the information supplied by one person can concern another**.

The Guide therefore remains focused on information rather than expanding into family governance or behavioural guidance.

Representation Is Not the Same as Understanding

There is another human-centred distinction to preserve.

The presence of extensive information about another person does not mean that an AI system possesses a complete understanding of that individual.

The information may represent only one side of a relationship.

It may describe isolated events.

It may contain the user's interpretations.

It may omit context known only to the other person.

Even documents that appear factual can represent only particular aspects of someone's life or activities.

An informational picture can therefore become detailed while remaining partial.

This matters because conversational fluency can make interpretations sound coherent. A response based on a detailed description may appear to capture the other person's circumstances convincingly.

But the quality of the language does not alter the limitations of the informational foundation.

A person represented through someone else's account remains represented through **someone else's account**.

The Information Relationship Is Not Always Two-Sided

The idea of “your data and AI” can initially suggest two parties:

you and **the AI**.

By this point in the Guide, that model is no longer sufficient.

The user may provide information about themselves.

Their interaction may reveal additional information.

Further information may be inferred.

Some information may persist and contribute to personalisation.

Separate details may combine into a larger picture.

And within that picture may be people who never interacted with the AI at all.

The information relationship can therefore become multi-person even when the interface remains one-to-one.

This is one of the most important consequences of recognising information as relational rather than merely individual.

It also prepares the ground for a broader question.

As information is deliberately supplied, indirectly revealed, inferred, accumulated and connected across people and contexts, users may begin to describe the resulting experience in simple human terms:

“The AI knows a lot about me.”

But what does that statement actually mean?

Does an AI “know” what a person explicitly told it?

Does it “know” an inference?

Does information available only in a current conversation constitute knowledge?

What if information is outdated or wrong?

What if the system can produce a highly personalised response while possessing only a partial informational picture?

The language of knowing can collapse several very different informational states into a single human concept.

Understanding those differences is the next step.

The question is deceptively simple:

What does AI actually mean when we say it “knows” something about us?

Chapter 12 — What Does AI Know About You?

People increasingly describe their interactions with artificial intelligence using the language of knowledge.

“The AI knows what I am working on.”

“It knows how I like things explained.”

“It already knows my background.”

“It seems to know what I am going to ask.”

Such expressions are understandable. An AI system that can respond to previous context, reflect stated preferences or produce answers apparently tailored to an individual's circumstances can create a strong impression of being known.

But the word *know* carries meanings developed through human experience.

When one person says that they know another, the statement can imply memory, familiarity, understanding, shared experience and awareness of the other person's circumstances. Human knowledge of another person exists within relationships and is continually interpreted through social context.

An AI system does not know a person in this human sense.

Yet simply saying that AI does not “really know” people would also leave an important part of the information relationship unexplained. AI systems can work with information about individuals. They can respond to contextual information, recognise relationships within available material and generate inferences from information that has been supplied. Depending on the system and environment, some information may also contribute to continuity or personalisation.

The more useful question is therefore not whether AI knows a person in an abstract sense.

It is:

What informational state are we describing when we say that AI “knows” something about us?

Knowing Can Mean Several Different Things

Consider a person who tells an AI system:

“I work in architecture.”

The information has been explicitly provided. If the system subsequently responds to a question using architectural examples, the user might reasonably say, “It knows I am an architect.”

But several different things may sit behind that statement.

The information might simply be available within the current conversation. It might form part of a broader contextual environment supported by the particular system. The response may also be drawing upon an inference from what the person said rather than relying upon a directly stated fact.

These informational states are different even though the resulting experience may feel similar.

This is why the language of knowing can be imprecise.

It can collapse **information supplied, information available, information retained, information inferred and information used** into one apparently simple idea.

Understanding the information relationship requires separating them again.

Information Supplied Is the Clearest Case

The most straightforward form of apparent knowledge is information that a person has explicitly provided.

A user states their occupation, describes a project, identifies their preferred response style or explains their circumstances. If that information remains available within the interaction, the AI can potentially use it when responding.

Nothing mysterious is required.

The system's apparent knowledge may simply reflect information present in its current context.

This distinction matters because people can attribute more understanding to a response than the informational circumstances justify.

An AI system that refers appropriately to a person's occupation may appear to possess a broader understanding of their professional life. In reality, the relevant informational foundation may consist of only a few statements.

The response can be highly coherent while the underlying picture remains narrow.

Fluency can therefore make limited information appear more comprehensive than it is.

Available Information Is Not Necessarily Permanent Information

A second distinction concerns availability.

Information can be available to an AI system for a particular interaction without necessarily having continuing relevance indefinitely.

Within an extended conversation, earlier information may remain part of the context and influence later responses. The user experiences continuity and may interpret that continuity as memory.

But, as Chapter 5 examined, persistence can take different forms. Information visible or usable in one context does not establish what will remain relevant in another.

The statement “AI knows this about me” can therefore obscure a temporal question:

For how long, and in what context, is the information available?

Something may function like knowledge within the current exchange while having no relevance to a later interaction.

Conversely, some environments may support forms of continuity that extend beyond one conversation.

The experience of being known can appear similar even though the informational architecture differs.

Retention and Use Are Not the Same Thing

Another distinction concerns the difference between information existing somewhere within an information environment and information actively shaping a response.

People can easily assume that if information has been retained in some form, the AI must be using it whenever they interact.

That does not necessarily follow.

Information can potentially exist without being relevant to a particular response. Similarly, a response may reflect information available in the immediate conversation without requiring a broader continuing record.

The concepts of **existence**, **availability** and **use** should therefore remain separate.

This is part of the visibility problem established earlier in the Guide. From the user's perspective, the principal evidence is the response itself. If the AI refers to a fact, the person knows that the information was relevant somehow. If it does not refer to the fact, the user cannot necessarily conclude that the information no longer exists anywhere within the wider environment.

The visible output does not provide a complete map of informational state.

Inference Can Feel Like Knowledge

The impression of being known becomes stronger when an AI system appears to recognise something the person never explicitly stated.

Suppose a user has been discussing several complex projects and repeatedly asking for highly technical explanations. The AI begins responding at an advanced level without being told explicitly, “I have substantial expertise in this field.”

The user may experience this as recognition:

“It knows I am experienced.”

Yet the informational basis may be an inference drawn from language, subject matter and interaction patterns.

As Chapter 4 established, inference differs from direct information.

The person did not supply the conclusion.

The available information supported it.

This distinction is particularly important because inferences can be wrong.

A researcher may ask expert-level questions about a field they are studying but do not practise professionally. A student may paste technically sophisticated material written by somebody else. A person may deliberately ask questions from a perspective different from their own.

An AI response can therefore appear to “know” something about a person when it is actually operating from an interpretation that may or may not accurately describe them.

Wrong Information Can Still Shape an Interaction

The language of knowing becomes still more problematic when the underlying information is inaccurate.

Suppose a person once states that they live in one country and later relocates. If the earlier information continues to influence subsequent interaction, the system may produce responses consistent with a circumstance that is no longer true.

Similarly, an inference made from an earlier conversation may have been incorrect from the beginning.

In both situations, the AI may appear to “know” something that is wrong.

This reveals an important difference between informational availability and truth.

An AI system can work coherently with inaccurate information.

A personalised response can be well constructed while resting on an outdated assumption.

A larger informational picture can contain errors while still appearing internally consistent.

The amount of information involved therefore does not determine whether the resulting representation of the person is accurate.

More information can create a richer picture.

It does not guarantee a truer one.

AI May Have Information Without Having a Person

Human beings usually understand personal knowledge holistically.

Knowing someone means more than possessing a collection of facts about them. It involves recognising that person as an individual whose circumstances, behaviour, intentions and relationships exist within a wider life.

AI information relationships can be much more fragmented.

A system may have access to information about a person's project without information about their broader professional circumstances. It may work with a stated preference without knowing whether that preference applies outside the present context. It may respond to detailed personal information while lacking many basic facts that a human acquaintance would know.

This creates a useful distinction between **having information about a person** and **having an understanding of the person**.

The first can occur computationally.

The second carries broader human meanings that should not automatically be transferred to AI.

A system may therefore appear remarkably knowledgeable in one moment and surprisingly unaware in the next. What seems inconsistent from a human relational perspective may make more sense when understood as differences in available informational context.

Personalisation Strengthens the Impression of Being Known

Chapter 6 examined how personalisation can create an accumulating informational picture.

The more effectively an AI interaction adapts to a person, the stronger the impression of knowledge may become.

A system that repeatedly provides responses in a preferred format can appear familiar with the user. One that works coherently with a long-running project can appear to understand its history. One that reflects recurring contextual information can feel increasingly individualised.

This experience is informationally significant even if the language of “knowing” remains imperfect.

People respond not only to technical realities but to how systems behave.

If AI behaves as though it recognises the individual, that behaviour can shape expectations about what information the system possesses.

A user may begin to assume that because the AI appears to know one thing, it also knows another.

That assumption may not be justified.

Personalisation can therefore expand both the informational relationship and the human perception of that relationship.

AI Can Also Appear to Know Other People

Chapter 11 demonstrated that information within an AI interaction can concern people other than the user.

The same problem of language applies here.

A person may provide extensive information about a colleague, relative or customer and then receive a detailed response concerning that individual. The AI may appear to “know” the other person.

But its informational picture may be based entirely upon material supplied by the user.

It may contain only one perspective.

It may include inference.

The other person may never have interacted with the system.

The resulting response can nevertheless sound coherent and specific.

This demonstrates why fluency should not be confused with completeness of informational understanding.

An AI can produce a detailed account from a partial informational foundation.

The Question “What Does It Know?” Has No Single Answer

By this stage, it should be clear why the apparently simple question “What does AI know about me?” cannot always be answered with a single category.

The answer may involve several different questions:

What information have I explicitly provided?

What information is available in the present interaction?

What information may have continuing relevance?

What has been inferred rather than stated?

What information is actually influencing the current response?

What information may be outdated or inaccurate?

What information concerns me but was supplied by somebody else?

These distinctions describe different parts of the information relationship.

The ordinary word *know* tends to merge them.

That does not mean people must stop using the word. Everyday language does not need to become technical every time someone interacts with AI.

But understanding requires recognising what the shorthand can conceal.

From “Knowing Me” to Data-Dependent Interaction

The question of what AI knows about an individual leads to a broader development.

AI interaction can become more useful when systems have relevant context.

Generic assistance requires relatively little information about the person. More contextual assistance can depend upon richer information. Personalised interaction can require informational differentiation. Continuing assistance can depend upon some form of informational continuity.

As these capabilities develop, the role of information can become increasingly central to the experience.

The significant change is not that AI suddenly begins to “know people” in a human sense.

It is that increasingly capable forms of interaction may depend upon increasingly rich informational context.

This moves the Guide beyond the question of what information exists about one individual and towards the structural relationship between **information and the development of AI interaction itself**.

The information relationship is not merely a side effect of using AI.

In many forms of contextual and personalised interaction, information is part of what makes the capability possible.

The next chapter therefore widens the perspective.

Instead of asking only what AI knows about an individual, it asks what happens as AI interaction becomes increasingly dependent upon information.

That is the broader development at the centre of **the growing data dependence of AI interaction.**

Chapter 13 — The Growing Data Dependence of AI Interaction

Artificial intelligence can provide useful assistance with very little information about the person using it.

A generic question can receive a generic explanation. A passage of text can be summarised without the system knowing who submitted it. A mathematical concept can be explained without knowing the learner's occupation, location or personal circumstances.

But many emerging forms of AI interaction move in another direction.

The more an AI system is expected to understand context, maintain continuity, adapt to preferences or provide individually relevant assistance, the more important information can become to the interaction. A generic system can respond to the question in front of it. A more contextual system may need to work not only with the question but with information surrounding the question.

This creates a broader development in the relationship between people and AI: **some forms of increasingly contextual and personalised AI interaction are also increasingly data-dependent.**

The word *dependent* requires care. It does not mean that every future AI capability must require large quantities of personal information, nor that all AI systems are developing in the same direction. Technical approaches, system designs and purposes differ. Some tasks can remain largely impersonal.

The important observation is more specific. Where the objective is for AI to become increasingly responsive to an individual's particular circumstances, information about those circumstances becomes part of what makes that responsiveness possible.

From Generic Capability to Contextual Capability

Consider a simple request:

“Explain how mortgages work.”

An AI system can answer without knowing anything about the person asking.

Now consider:

“Explain how the mortgage options I am looking at differ in relation to my circumstances.”

The second task requires context.

The system may need information about the options being considered and aspects of the person's situation relevant to explaining the differences. The more individually specific the requested explanation becomes, the more informational context may be required.

The same principle appears in many domains.

A generic explanation of how to prepare for a job interview requires little personal information. Assistance tailored to a person's experience, prospective role and previous interview difficulties requires more.

A generic explanation of project planning can be produced with minimal context. Assistance with a continuing project requires information about that project, its history and current circumstances.

These examples do not establish what decisions people should make from AI assistance. Decision-making remains a distinct subject within the wider AIGuide series.

They illustrate an informational principle:

greater contextual relevance often requires greater contextual information.

Information Can Become Part of the Capability

This principle changes how data can be understood within AI interaction.

Information is sometimes treated as something external to the capability: the AI has its capabilities, and the user simply supplies data for those capabilities to process.

In contextual interaction, the distinction becomes less clear.

The ability to respond appropriately to a person's particular circumstances can itself depend upon having relevant information about those circumstances.

Without the context, the system may still possess substantial general capability. What it lacks is the informational basis for applying that capability in a specifically contextual way.

This means information can become part of the practical conditions under which certain forms of AI usefulness emerge.

The AI capability and the information remain conceptually distinct, but the experience available to the user can depend upon their combination.

Continuity Creates Further Information Dependence

The same relationship appears with continuity.

If every AI conversation begins entirely afresh, the user may need to restate relevant background repeatedly.

A continuing interaction can reduce that repetition.

But continuity requires some connection between the present and the past.

As Chapter 5 examined, persistence can take different forms and should not be reduced to the simple claim that an AI “remembers.” Nevertheless, wherever earlier context is relevant to a later interaction, information has acquired a temporal role.

The system's usefulness is no longer based only upon what the user says now.

It may also depend upon information that establishes what came before.

This becomes especially significant in longer-running activities.

A person working on an extended project may value an AI interaction that can operate within an established context. Someone developing knowledge of a subject over time may benefit when the interaction reflects what has already been covered. Repeatedly reconstructing the complete background can reduce the value of continuity.

The desire for continuing assistance can therefore create a corresponding informational requirement.

Continuity has an information cost in the neutral sense that continuity requires information about what is being continued.

That is not a judgement about whether the cost is acceptable. It describes the structural relationship between the capability and the information that enables it.

Personalisation Deepens the Relationship

Personalisation adds another layer.

A generic AI response can be useful to many people in approximately the same form. A personalised response attempts to reflect differences among individuals or contexts.

As Chapter 6 established, such differentiation has an informational foundation.

The system needs some basis for adapting.

That basis might be very modest: a preference stated within the current conversation. Or it may involve a richer informational context, depending upon the system and purpose.

As personalisation becomes more sophisticated, the potential range of relevant information can broaden.

Preferences may matter.

Previous context may matter.

Ongoing activities may matter.

Information about the person's current circumstances may matter.

Patterns of interaction may matter.

This does not mean that all available information is necessary or that more data automatically produces better personalisation. Information can be irrelevant, outdated, incomplete or misleading.

The relationship is therefore not simply:

more information = better AI.

A more accurate formulation is:

some forms of more individually relevant AI require relevant information capable of supporting that relevance.

The distinction matters because quantity and usefulness are not interchangeable.

Richer Information Can Also Produce Richer Errors

Data dependence introduces an important limitation.

If contextual AI depends upon information, then problems in that information can affect the resulting interaction.

Outdated context can support outdated responses.

Incorrect information can create inappropriate assumptions.

An inaccurate inference can become part of a broader picture.

Information that was relevant in one setting may be carried conceptually into another where it is less appropriate.

This means that increased contextual richness does not guarantee increased contextual accuracy.

Indeed, a highly personalised response can sometimes appear more convincing precisely because it reflects details associated with the individual. If some of those details are wrong, the response may still sound coherent.

The information relationship therefore becomes more consequential as information plays a greater role in shaping the interaction.

Again, this does not determine whether the output should be trusted. The separate question of trust belongs primarily to AIGuide 7, *Knowing When to Trust AI*.

For AIGuide 12, the relevant point is informational: **when an AI experience depends more heavily on context, the quality and relevance of that context become more significant to the character of the interaction.**

Data Dependence Can Expand Beyond the Individual

Contextual AI may also depend upon information that is not solely about the person using it.

A workplace project involves colleagues and organisational information. A family situation involves other family members. A document may contain information about customers or clients. A continuing activity may be understandable only through relationships among several people.

Chapter 11 established that AI information relationships can therefore become multi-person even when only one person operates the interface.

Growing data dependence can deepen this effect.

If contextual usefulness depends upon richer descriptions of real-world situations, and real-world situations involve other people, then the informational foundations of AI assistance may extend beyond the individual user.

This reinforces why “personalisation” should not be understood only as the construction of a picture of one isolated person.

The relevant informational context may include projects, relationships, organisations, places and events.

AI interaction can become more context-dependent at the same time that the context itself becomes more socially interconnected.

Multimodal Interaction Broadens What Counts as Context

The growing information dependence of AI is also visible in the range of material systems can work with.

Interaction is no longer necessarily confined to typed text.

Documents, images, audio and other forms of material can become part of the context within which assistance is provided. Different forms of information can potentially contribute different dimensions to an interaction.

This broadens the meaning of data dependence.

A contextual AI experience may not depend solely upon a larger quantity of text. It may depend upon a richer variety of information.

An image can provide context difficult to communicate verbally.

A document can introduce detailed material without requiring the user to reproduce it manually.

Audio can make spoken material available for processing.

The resulting interaction can become more capable because the system has access to forms of context better suited to the task.

But the principle established in Chapter 2 remains relevant: the informational content of these materials can extend beyond the immediate purpose for which they were supplied.

Richer context can create richer capability while simultaneously creating a richer information relationship.

Increasing Capability Does Not Require One Model of Information Use

It would be misleading to treat data dependence as a prediction that all AI systems will evolve towards one model in which ever-increasing amounts of personal information are accumulated.

There are many possible approaches.

Different systems may minimise continuing personal context, rely heavily on immediate information, support selective continuity, allow different forms of user control or operate within specialised environments where particular information is relevant only for defined purposes.

Technical development can also change how much information is required to accomplish particular tasks.

The relationship between AI capability and information is therefore not fixed.

This Guide is concerned with the direction of the conceptual trade-off rather than prescribing a particular architecture.

Where people expect AI to become more aware of their circumstances, there must be some informational basis for that awareness.

How that basis is created, limited, structured or managed can vary substantially.

Expectations May Become Data-Dependent Too

As people experience more contextual AI, their expectations can change.

A generic answer that once seemed impressive may later feel unnecessarily impersonal. Users accustomed to continuity may find repeated explanations frustrating. Someone who has

experienced effective personalisation may increasingly expect systems to recognise relevant context automatically.

This creates a human dimension to data dependence.

People may come to value capabilities that themselves depend upon richer information while simultaneously wanting clearer understanding of the information relationship that makes those capabilities possible.

These preferences are not necessarily contradictory.

A person can value personalisation and still want to understand its informational foundation.

They can appreciate continuity while wondering what persists.

They can value context-sensitive assistance while remaining uncertain about what information contributes to it.

The challenge is therefore not captured by a simple opposition between usefulness and privacy.

The relationship is more complex.

Usefulness itself can sometimes be information-dependent.

The Information Relationship Is Becoming Part of the AI Experience

This chapter marks an important widening of perspective.

Earlier chapters examined individual dimensions of information: what people provide, what they reveal, what may be inferred, what persists, what accumulates and what can involve other people.

Taken together, these dimensions point towards a larger development.

For increasingly contextual forms of AI interaction, the information relationship is not merely something occurring in the background.

It can become part of the experience itself.

The AI appears more relevant because context is available.

It appears more continuous because previous information remains relevant in some form.

It appears more personalised because information differentiates one interaction from another.

It may appear to “know” the user because information about the person is reflected in its responses.

Information and interaction consequently become increasingly intertwined.

A Relationship That Requires Understanding

The growing data dependence of AI interaction does not produce a simple conclusion.

It does not establish that more contextual AI is necessarily better.

It does not establish that generic AI is preferable.

It does not tell individuals what information they should provide.

And it does not determine how organisations or providers should design their systems.

Those would be prescriptive conclusions beyond the purpose of this Guide.

What can be observed is a developing relationship.

As AI interaction becomes more contextual, continuous and personalised, relevant information can become increasingly important to what the system is able to provide.

That creates both capability and complexity.

The capability lies in AI becoming more responsive to particular circumstances.

The complexity lies in the increasingly rich information relationship that can support that responsiveness.

And much of that relationship may remain difficult for the ordinary user to see as a whole.

This brings the Guide towards its final synthesis.

The central challenge is no longer simply identifying individual categories of information. It is understanding how people live with an information relationship that can be useful, dynamic, cumulative, partially inferred, socially interconnected and not always completely visible.

The final chapter therefore returns to the question with which the Guide began.

When we interact with AI, what information relationship are we creating—and how well do we understand it?

That is the challenge of **living with an information relationship we cannot fully see.**

Chapter 14 — Living with an Information Relationship We Cannot Fully See

The information relationship between people and artificial intelligence can begin almost invisibly.

A person asks a question.

They provide some context.

The AI responds.

Nothing about the exchange necessarily feels like the creation of an information relationship. The person's attention is usually directed towards the purpose of the interaction: understanding something, analysing material, developing an idea, preparing a document or obtaining assistance with a task.

Yet across the preceding chapters, a wider picture has emerged.

Information can be deliberately supplied. It can also be revealed indirectly through language, context and patterns of interaction. Information can provide the basis for inference. It can have significance beyond the moment in which it was introduced. It can contribute to personalisation and continuity. Separate pieces can combine into a larger picture. Information provided by one person can concern others. And as AI interaction becomes increasingly contextual, relevant information can become part of what makes that interaction useful.

The resulting relationship is neither inherently beneficial nor inherently harmful.

It is, however, increasingly important to understand.

The difficulty is that the relationship cannot always be seen in its entirety.

We See the Conversation More Easily Than the Information Relationship

The visible AI interaction is usually straightforward.

There is an interface.

There is information entered or uploaded.

There is an output.

This visible exchange provides the user with an intuitive mental model of what has happened: “I gave the AI this, and it gave me that.”

That model is not wrong.

It is incomplete.

The information deliberately entered into a system is only one dimension of the relationship examined throughout this Guide.

A person's words can reveal context beyond their explicit statements. Repeated interactions can create patterns. Available information can support inferences. Earlier context can sometimes remain relevant to later interaction. Personalisation can make information part of how an AI system differentiates one context from another.

None of these dimensions is necessarily obvious from the interface.

This creates the visibility gap examined in Chapter 9: the difference between what users can readily observe and the broader information relationship in which they may be participating.

The existence of such a gap does not mean that everything beyond the interface is hidden or problematic.

It means that **using an AI system and completely understanding its information environment are different activities.**

Complete Visibility May Be an Unrealistic Standard

There is a temptation, once informational complexity is recognised, to imagine that complete visibility would resolve the problem.

If every process could be seen, perhaps the relationship would become fully understandable.

But complex technologies do not generally operate that way.

People use systems every day without seeing their complete technical operation. They do not observe every process involved in a financial transaction, mobile communication or online service. The fact that some processes remain outside immediate view does not make meaningful participation impossible.

AI is similarly complex.

Complete visibility into every technical process would not necessarily produce complete human understanding. It could instead overwhelm ordinary users with information having little relevance to the decisions or tasks immediately before them.

The challenge is therefore more subtle than making everything visible.

It is recognising that **partial visibility requires humility about what can reasonably be assumed from what we see.**

The visible interaction tells us something important.

It does not necessarily tell us everything.

More Information Does Not Necessarily Mean More Understanding

The same distinction applies to the amount of information surrounding a person.

Across repeated interactions, information can accumulate.

At first glance, this might suggest a simple progression:

more information leads to a more complete picture, which leads to better understanding.

The reality is less straightforward.

Information can be outdated.

It can be incomplete.

It can describe temporary circumstances.

It can have been provided hypothetically.

It can reflect one context but not another.

Information supplied by another person may represent only that person's perspective.

Inferences may be plausible without being correct.

A larger informational picture can therefore become more detailed without becoming proportionately more accurate.

This is one of the most important distinctions in understanding AI and data.

Information richness is not the same as human understanding.

An AI system may work with a considerable amount of information concerning an individual while lacking many aspects of context that a human acquaintance might regard as fundamental.

Conversely, it may possess very little information about a person and still provide highly useful assistance for a particular task.

The informational relationship should therefore not be measured only by quantity.

The Relationship Can Be Useful Because It Contains Information

It would also be incomplete to examine the information relationship only through questions of exposure or uncertainty.

Information enables capability.

An AI system can provide a generic explanation with minimal knowledge of the user. But if the person wants assistance tailored to their particular circumstances, relevant information can improve contextual usefulness.

Continuity can reduce repetition.

Personalisation can make interaction more individually relevant.

Documents, images and other materials can allow AI to work directly with information that would otherwise be difficult to describe.

Context can make complex tasks easier to pursue across extended conversations.

These benefits are part of the same information relationship.

This is why a simple framing in which “less information is good” and “more information is bad” does not adequately describe the subject.

Nor does the opposite framing—that greater personalisation and information availability are automatically desirable.

The relationship contains different considerations that can coexist.

Information can create usefulness while also creating informational complexity.

Context can improve relevance while remaining incomplete.

Persistence can support continuity while preserving information whose relevance changes.

Inference can make interaction more responsive while producing conclusions that may be wrong.

The same informational capability can therefore have more than one dimension.

Human Expectations Are Part of the Relationship

The information relationship is not defined only by technical systems.

It is also shaped by what people believe those systems are doing.

Conversational AI can feel private because the exchange appears to involve only one person and one system.

Personalisation can create the impression of being recognised.

Continuity can feel like memory.

A well-contextualised response can feel like understanding.

An accurate inference can feel like knowledge.

These perceptions matter because they influence how people experience AI.

But the human concepts underlying them—privacy, memory, recognition, understanding and knowledge—carry meanings developed through relationships among people.

Computational information processes are different.

Recognising that difference does not require stripping conversational AI of all familiar language. People will naturally continue to say that an AI “remembers,” “knows” or “understands.”

The important point is to recognise the shorthand.

Behind apparent memory may be some form of informational continuity.

Behind apparent knowledge may be available information.

Behind apparent recognition may be personalisation.

Behind apparent understanding may be contextual processing and inference.

The human experience is real.

The mechanism should not automatically be assumed to be human-like.

Other People Are Part of Our Information World

The information relationship also cannot be understood purely in individual terms.

People live in relationships.

Their messages contain other people's words. Their photographs contain other people's images. Their documents refer to colleagues, customers, students, family members and organisations. Their descriptions of everyday life inevitably include others.

As Chapter 11 established, the person providing information and the person represented by that information are not necessarily the same.

This means that an AI interaction can have an informational reach extending beyond its visible participant.

Someone can become part of an informational picture without personally interacting with the AI system.

Again, this observation does not determine what should be done in any particular situation. Professional, legal, organisational and personal circumstances differ too widely for a universal rule, and providing such advice would exceed the scope of this Guide.

It does, however, change the conceptual picture.

The information relationship between humans and AI is not always simply **my information and my AI interaction**.

It can involve information embedded within the social world around us.

The Relationship Will Not Remain Static

The significance of these questions is likely to evolve because AI interaction itself continues to develop.

Systems can increasingly work with different forms of material. Contextual capabilities can become richer. AI may be incorporated into more services and activities. The boundaries between isolated tasks and continuing assistance can shift.

This does not justify certainty about the future.

Technologies, business models, regulation, social expectations and technical architectures can all develop in different directions. Some environments may support richer personalisation. Others may deliberately limit informational continuity. New approaches may alter how information is handled or how much context particular capabilities require.

The durable issue is therefore not any one current technical arrangement.

It is the underlying relationship between **information and increasingly capable AI interaction**.

As long as AI systems use information to interpret context and provide assistance, understanding that relationship will remain relevant even as particular technologies change.

Understanding Without Prescription

A natural response to the issues examined in this Guide is to ask:

What should I tell AI?

What should I withhold?

Which settings should I choose?

Which system should I trust with particular information?

Those are understandable questions.

They are not questions this Guide attempts to answer.

The circumstances surrounding information can differ dramatically. A generic public interaction, an organisational AI environment and a specialised professional system may involve very different considerations. Legal and regulatory requirements vary internationally. Individuals also differ in their purposes, circumstances and judgements.

Providing universal behavioural instructions would therefore oversimplify the very complexity the Guide has sought to make visible.

AISF's purpose here is more foundational.

Before a person can exercise judgement, they need a sufficiently clear understanding of what they are judging.

That means recognising that an AI interaction has an information dimension.

It means recognising that information extends beyond obvious identifiers.

It means distinguishing what is directly provided from what is indirectly revealed or inferred.

It means understanding that persistence, personalisation and aggregation can change the informational context.

It means recognising that other people's information may also be involved.

And it means distinguishing the familiar experience of conversation from the more complex information environment within which that conversation occurs.

Understanding does not dictate the decision.

It improves the basis on which human judgement can operate.

Living with Partial Knowledge

Perhaps the most realistic conclusion is that people will often interact with AI without possessing a complete picture of the information relationship.

That condition is not necessarily unique to AI, nor is it necessarily avoidable.

Modern technological environments are complex.

The relevant question is therefore not whether every person can achieve perfect informational awareness.

It is whether people can develop a sufficiently clear conceptual understanding to recognise the dimensions that exist even when they cannot see every process directly.

The distinction matters.

Someone who assumes that an AI interaction consists only of the words visible in a current prompt has one mental model.

Someone who understands that information can also be contextual, inferred, persistent, aggregated, personalised and relational has another.

The second person does not necessarily know every technical detail.

But they can understand the nature of the questions that exist.

That is a meaningful difference.

Returning to the Central Question

This Guide began with a question:

When we interact with AI, what information relationship are we creating—and how well do we understand it?

There is no single answer applicable to every system, every person or every interaction.

Sometimes the relationship may be minimal: a generic question containing little information about the person asking it.

Sometimes it may be rich: an extended interaction involving personal circumstances, uploaded material, continuing context and information about other people.

Sometimes the user may have a relatively clear understanding of the information involved.

At other times, the informational significance of the interaction may be much harder to see.

The relationship exists along a spectrum rather than as a single condition.

What remains constant is the importance of recognising that the relationship exists at all.

AI does not only provide information to people.

People also bring information into AI interaction.

That information can describe them, reveal aspects of their circumstances, contribute to inference, extend through time, support personalisation, combine into larger pictures and include the lives of others.

As AI becomes more capable of working with context, this information relationship may become increasingly central to the usefulness people experience.

The challenge is therefore not to eliminate information from AI interaction.

Nor is it to accept every information relationship without question.

It is to **see the relationship clearly enough for human judgement to remain meaningful.**

That returns AIGuide 12 to the human-centred principle underlying the wider AISF AI Guidance series.

Technology can process information.

AI can identify patterns, work with context and generate responses.

Systems can become increasingly capable of adapting to the information available to them.

But understanding what that relationship means remains a human task.

The information may be complex.

The environment may be only partly visible.

The technology may continue to change.

Human judgement therefore does not become less important as AI becomes more capable.

It becomes part of how people make sense of the information relationship they are creating.

Key Ideas from This Guide

Several ideas run through AIGuide 12.

Every AI interaction has an information dimension.

Attention naturally falls on what AI gives the person, but understanding the interaction also requires considering what information the person brings into it.

Information is broader than formal data.

Words, documents, images, audio, context, patterns of interaction and information concerning other people can all contribute to the informational environment.

What we deliberately provide is only part of the picture.

Language, context and repeated interactions can reveal information beyond what a person consciously regards as a disclosure.

Information and inference are different.

Available information can support conclusions that were never explicitly supplied. Those conclusions may be plausible without necessarily being correct.

Information can have significance beyond the immediate interaction.

The end of a visible conversation does not necessarily mean the end of the informational significance of everything introduced within it.

Personalisation depends upon information.

Context can make AI interaction more relevant and useful. The same process can also make the informational foundations of the interaction more significant.

Conversation and confidentiality are not equivalent.

An interaction may feel direct, individual and conversational without thereby possessing the characteristics associated with confidential human or professional relationships.

Participation and complete understanding are different.

A person can successfully use an AI-enabled service without possessing a complete picture of the information environment in which the interaction occurs.

There can be a visibility gap.

Users readily see some aspects of an interaction—their prompts, uploads and outputs—while other informational dimensions may be less immediately apparent.

Small pieces can create a larger picture.

Information that appears unremarkable when considered separately can acquire different significance when combined with other information.

Your AI interaction may contain information about other people.

Documents, messages, photographs, workplace materials and descriptions of everyday life can bring people who are not themselves interacting with the AI into the informational picture.

AI “knowing” a person is not the same as human knowledge of another person.

What appears as recognition or understanding may instead involve available information, contextual processing, personalisation and inference.

More information does not necessarily mean more understanding.

Information can be incomplete, outdated, hypothetical or context-dependent. A richer informational picture is not automatically a more accurate human picture.

Information enables capability.

The information relationship should not be understood solely through exposure or uncertainty. Context, continuity and relevant information can materially increase the usefulness of AI interaction.

And finally:

Complete visibility may be unrealistic, but meaningful understanding remains possible.

A person does not necessarily need to understand every technical process to recognise that information can be directly supplied, indirectly revealed, inferred, persistent, aggregated, personalised and relational.

That conceptual understanding matters because:

Human judgement does not become less important as AI becomes more capable. It becomes part of how people make sense of the information relationship they are creating.

AI Guidance — Continuing the Exploration

AIGuide 12 examines one particular dimension of artificial intelligence: **the information relationship created when people interact with AI.**

Some neighbouring subjects are deliberately bounded to other Guides.

AIGuide 11 — AI in the Family addresses the broader questions surrounding AI within family relationships, including roles, responsibilities, household boundaries and the particular implications of AI in family life. AIGuide 12 uses family circumstances where they illuminate information involving other people, but does not expand that discussion into the wider family-AI relationship.

AIGuide 7 — Knowing When to Trust AI addresses the separate question of whether AI or a particular AI output should be trusted. AIGuide 12 remains focused on the information upon which increasingly contextual AI interaction may depend and the consequences of that informational dependence.

These relationships do not make AIGuide 12 dependent upon the other Guides.

It is intended to stand independently while enabling readers who wish to explore neighbouring subjects to continue through the wider AI Guidance series.

Future Edition Updates & User Submissions

The information relationship between people and artificial intelligence will continue to develop.

AI systems may become more contextual, multimodal and personalised. Information-handling practices may change. Regulatory approaches may develop differently across jurisdictions. New technical architectures may alter how information is processed, retained, separated or made available to AI-enabled capabilities.

The terminology used by providers and users may also evolve.

AISF therefore treats AI Guidance as an evolving educational resource rather than a permanently fixed description of a changing technological environment.

Future editions may incorporate material developments where they meaningfully affect the concepts examined in this Guide.

AISF also welcomes observations from readers concerning emerging information practices, areas where additional explanation could improve understanding and subjects that may warrant consideration in future AI Guidance publications.

Reader submissions may inform future consideration but do not constitute endorsement, acceptance or a commitment to incorporate submitted material.

AISF retains independent editorial responsibility for all published content.

Acknowledgements

This Guide was developed through a structured human–AI collaborative process under the editorial direction and governance of AI Sourced Facts (AISF).

Artificial intelligence was used as an enabling capability in research, structured examination, drafting, challenge and refinement.

Human oversight remained integral throughout development, including determination of purpose, scope, architecture, editorial standards, substantive acceptance and final publication.

AISF also acknowledges the wider community of researchers, educators, practitioners, institutions, technology developers, regulators and users whose continuing examination of artificial intelligence and information contributes to broader public understanding.

Responsibility for the final content and editorial position of this Guide remains with AISF.

AISF Universal Disclaimer

This publication is provided for **general educational and informational purposes only**.

It is intended to support understanding of issues associated with artificial intelligence and information and should not be interpreted as professional, legal, privacy, cybersecurity, financial, medical, regulatory, technical or other specialised advice.

The Guide does not recommend or endorse any particular artificial intelligence system, provider, product, platform, service, information practice, privacy setting or course of action.

AI technologies, technical architectures, services, information practices and regulatory requirements can change rapidly and may differ substantially between systems, providers, jurisdictions and operating environments.

Examples concerning information, inference, persistence, personalisation, confidentiality, consent, aggregation and AI “knowledge” are intended to explain concepts. They should not be interpreted as statements that every AI system performs, retains, processes or uses information in the same manner.

Readers should consult current provider information, applicable policies, authoritative regulatory sources and suitably qualified professionals where particular circumstances require technical, legal, privacy, cybersecurity or other specialist guidance.

AISF seeks to maintain neutrality, clarity and appropriate recognition of uncertainty. No publication can eliminate all possibility of error, omission, subsequent change or differing interpretation.

Readers remain responsible for the use they make of the information contained in this Guide.

About AI Sourced Facts (AISF) Pte. Ltd.

AISF is a Singapore-headquartered institution dedicated to structured reasoning, responsible AI navigation, and governance-informed adoption of artificial intelligence systems.

AISF operates with a capability-first, vendor-neutral posture. Its publications do not rank platforms, endorse providers, or promote specific technologies. Instead, AISF develops structured frameworks that help individuals, professionals, and institutions reason clearly before integrating AI into operational, strategic, or educational environments.

AISF's work spans whitepapers, applied insight books, education instruments, governance architectures, and structured research initiatives. These outputs are informed by cross-system AI research methodologies and reflect globally observed usage patterns at the time of publication. Human accountability remains central across all AISF frameworks.

AISF does not provide regulatory, legal, financial, investment, or compliance advice. Its publications are designed to support structured thinking, proportionate governance, and disciplined evaluation of AI capabilities prior to deployment or reliance.

As artificial intelligence systems continue to evolve, AISF's focus remains constant: clarity before integration, governance proportionate to capability, and long-term institutional resilience in the age of AI.

AI Sourced Facts (AISF)

Truth, Clarity & Guidance in the Age of AI.

Publication & Governance Information

Title: Your Data and AI

Subtitle: Understanding What We Give AI—and What May Follow

Series: AI Guidance

Guide Number: AIGuide 12

Edition: 2026 Edition

Publisher: AI Sourced Facts (AISF)

Copyright © 2026 AI Sourced Facts (AISF). All rights reserved.

This publication forms part of the AISF AI Guidance series.

The content may be updated in future editions to reflect material developments in artificial intelligence, information practices and the wider environment in which people interact with AI.

Back Cover

Your Data and AI

You ask AI a question.

You upload a document.

You describe a situation.

You provide some context.

The AI responds.

The exchange may feel simple.

But what information relationship has been created?

AIGuide 12 — Your Data and AI examines the information dimension of everyday artificial intelligence interaction.

Across fourteen chapters, the Guide explores what people deliberately give AI, what interaction can reveal indirectly, inference, persistence, personalisation, confidentiality, consent, visibility, aggregation, information about other people and the growing dependence of contextual AI upon information.

It does not tell readers what they should disclose or which systems they should use.

Instead, it helps make a relationship that can be difficult to see more understandable.

Because AI does not only provide information to people.

People also bring information into AI interaction.

And as AI becomes increasingly capable of working with context, understanding that relationship becomes an increasingly important part of informed human judgement.

Part of the **AISF AI Guidance series** — structured, neutral and human-centred examinations of artificial intelligence and its place in everyday life.

AI Sourced Facts (AISF)

Truth, Clarity & Guidance in the Age of AI.

Website: www.aisourcedfacts.com

Quick Access: AISF.tech