



Living with AI-Generated Content

(AIGuide 10)

Understanding Origin, Authenticity and Interpretation in an
Increasingly Synthetic Information Environment

AI Sourced Facts (AISF) Pte. Ltd.

© AI Sourced Facts (AISF)
All rights reserved.

Edition: Version 1.0

Publication Date: August 2026

Preface

For much of everyday life, people have been able to make quiet assumptions about the content they encounter.

A photograph suggested that something had been placed before a camera. A recording suggested that someone had spoken. A video appeared to preserve some connection with an event that occurred. Written material, whatever its quality, was generally assumed to have originated through human authorship.

Those assumptions were never perfect.

Photographs could be manipulated. Recordings could be edited. Films could construct events that never happened. Written material could be anonymous, copied, fabricated or falsely attributed.

Artificial intelligence does not create the distinction between appearance and reality.

It changes its scale, accessibility and complexity.

Text can now be generated without a person composing every sentence. Images can depict people and places that were never photographed. Voices can be synthesised or transformed. Video can combine recorded and generated elements. Existing content can be altered in ways ranging from minor assistance to substantial transformation.

As these capabilities develop, the surface of content can reveal progressively less about the process that produced it. This is the foundational change with which AIGuide 10 begins.

The resulting challenge is not adequately captured by asking whether something is simply “real” or “fake.”

Human and machine contributions can coexist. Synthetic content can be openly fictional, educational or creative. Human-created content can be misleading. Generated material can accurately communicate information. Authenticity can concern source, authorship, representation, experience or production.

AIGuide 10 examines this changing informational environment.

Its central question is:

What changes when people can no longer assume how content was created simply from seeing, hearing or reading it?

Purpose of the Guide

The purpose of **AIGuide 10 — Living with AI-Generated Content** is to help readers understand how increasingly capable generation and modification technologies are changing the relationship between **content, origin, authenticity, context and human interpretation**.

The Guide examines:

the weakening relationship between appearance and origin;

the continuum between human-created and machine-generated content;

the different signals associated with text, images, voices and video;

the changing meaning of authenticity;

provenance and production history;

synthetic representations of people, voices, places and events;

generation and distribution at scale;

the blurring of information, expression and simulation;

the possibility that genuine content may itself become easier to doubt;

familiarity, repetition and apparent credibility;

shared reality within an increasingly synthetic information environment;

the importance of context;

interpretation under uncertainty; and

the continuing evolution of the information environment.

The Guide does not attempt to provide a technical manual for detecting AI-generated content.

It does not assume that synthetic content is inherently deceptive or harmful.

It does not equate human-created content with truth.

It does not attempt to establish universal legal or regulatory definitions of AI-generated material.

It does not provide detailed factual-verification methodology, a general framework for deciding when AI should be trusted, or the broader architecture of AI-assisted human decision-making. Those subjects occupy neighbouring territory elsewhere within the AI Guidance series. Its purpose is more fundamental:

to help people understand what it means to interpret content when appearance alone increasingly does not reveal the history beneath it.

Global Neutrality Statement

AI-generated content is a global phenomenon, but the ways societies interpret and respond to it are not necessarily uniform.

Expectations surrounding authorship, originality, creativity, disclosure, identity and acceptable technological assistance can vary across cultures, professions, institutions and jurisdictions.

Different countries may establish different regulatory approaches.

Different industries may develop their own conventions.

The significance attached to AI involvement may differ between journalism, education, advertising, entertainment, creative work, workplace communication and personal communication.

Technological capabilities may also develop unevenly across languages, regions and communities.

This Guide therefore does not assume that one definition, classification or social expectation surrounding AI-generated content applies universally.

Examples are illustrative.

The Guide is system-neutral and does not endorse or rank particular AI systems, generation technologies, platforms, detection services, provenance systems or regulatory approaches.

Its concern is the changing **human information environment**, not advocacy for or against particular technologies.

Editorial Standard

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

Accordingly, AIGuide 10 does not proceed from the assumption that AI-generated content is inherently false, deceptive, inferior or undesirable.

Nor does it assume that human-created content is inherently authentic, accurate or trustworthy.

The Guide avoids reducing the emerging environment to a binary contest between **human and machine, real and fake, or authentic and artificial**.

Instead, it examines relationships.

How was content produced?

What does it appear to represent?

What is known about where it came from?

What roles did humans and machines play?

What does authenticity mean in the particular context?

What remains uncertain?

These questions do not prescribe a conclusion.

They help expose dimensions of content that may no longer be visible from its finished surface.

The editorial objective is therefore not to tell readers how they should react to AI-generated content.

It is to make the environment in which they encounter it **more intelligible**.

How to Use This Guide

AIGuide 10 is designed as a progressive examination of an information environment in transition.

The opening chapters establish the central change: content increasingly may not reveal its origins, and the familiar binary between human-created and machine-generated material is giving way to a wider continuum of assistance, modification, transformation and generation.

The Guide then examines how this change manifests differently across text, images, voices and video before turning to authenticity and provenance.

Subsequent chapters consider synthetic representations of people, places and events; the significance of producing content at scale; and the increasingly blurred relationships among information, expression and simulation.

The later chapters examine broader consequences: genuine content becoming easier to doubt, repetition contributing to apparent credibility, pressures upon shared informational reality, and the significance of context.

The final chapters return explicitly to the human position: interpreting content under uncertainty and living within an information environment that will continue to change.

Readers may consult individual chapters independently.

However, the Guide is designed to be read progressively because its distinctions accumulate.

This is not a detection manual.

It is a framework for understanding **what changes when the surface of content no longer necessarily reveals its origin.**

Table of Contents

Preface

Purpose of the Guide

Global Neutrality Statement
Editorial Standard
How to Use This Guide
Version & Governance

Chapters

Chapter 1 — When Content No Longer Reveals Its Origins

Chapter 2 — From Human-Created to Machine-Generated: A Growing Continuum

Chapter 3 — Text, Images, Voices and Video: Different Forms, Different Signals

Chapter 4 — What Does “Authentic” Mean Now?

Chapter 5 — Provenance: Understanding Where Content Came From

Chapter 6 — Synthetic People, Voices, Places and Events

Chapter 7 — When Content Can Be Produced at Scale

Chapter 8 — The Blurring of Information, Expression and Simulation

Chapter 9 — When Genuine Content Can Also Be Doubted

Chapter 10 — Familiarity, Repetition and the Appearance of Credibility

Chapter 11 — Shared Reality in an Environment of Synthetic Content

Chapter 12 — Different Contexts, Different Consequences

Chapter 13 — Interpreting Content Under Uncertainty

Chapter 14 — Living with an Information Environment That Keeps Changing

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 10 — Living with AI-Generated Content

Understanding Origin, Authenticity and Interpretation in an Increasingly Synthetic Information Environment

Edition: 2026 Edition

Series: AISF AI Guidance

Published by: AI Sourced Facts (AISF)

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

The technologies capable of generating and modifying text, images, audio and video continue to develop. Social conventions, institutional practices, expectations surrounding disclosure and methods of communicating content provenance may develop alongside them.

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

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

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

Chapter 1 — When Content No Longer Reveals Its Origins

For much of everyday life, people have been accustomed to making quiet assumptions about the things they read, see and hear. A photograph suggested that a camera had captured something placed before it. A recording suggested that a voice had spoken the words being heard. A video appeared to preserve some visual connection with an event that had occurred. Written material, whatever its quality or accuracy, was generally understood to have originated through human authorship. These assumptions were never completely reliable, but they provided a familiar starting point for interpreting content.

Artificial intelligence is changing that starting point.

Text can now be generated without a person composing each sentence. Images can depict people, objects and places that were never photographed. Voices can be synthesised or modified in ways that resemble human speech. Video can show events, actions or expressions that did not occur in the form presented. Existing material can also be altered through AI in ways ranging from minor enhancement to substantial transformation. As these capabilities become more accessible and their outputs become more convincing, the visible or audible characteristics of content reveal less about how that content came into existence.

This does not mean that everything has suddenly become artificial, nor that people can no longer understand anything they encounter. It means something more specific: **the relationship between the appearance of content and its origin is becoming less obvious.**

That distinction is central to living with AI-generated content.

The assumptions behind ordinary interpretation

People have always interpreted information using more than the literal content placed before them. Context, source, presentation, experience and expectations all contribute to meaning. A photograph appearing in a newspaper, for example, is encountered differently from an illustration in a novel. A recorded voice message from a familiar person carries different assumptions from a fictional voice in an animated film. The same image may be interpreted differently depending on whether it appears in journalism, advertising, entertainment or satire.

Technology has complicated these assumptions before. Photographs could be staged or manipulated long before generative AI. Film editing could rearrange events. Audio recordings could be cut and combined. Computer-generated imagery created convincing scenes that had never existed outside a production environment. Digital editing tools made alteration progressively easier. Written material could always be anonymous, ghostwritten, copied, fabricated or falsely attributed.

AI-generated content therefore does not create the distinction between appearance and reality.

What changes is the extent to which content can be created or transformed rapidly, accessibly and across multiple forms without its method of creation necessarily being apparent from the finished result.

A realistic photograph-like image may now have no photographed scene behind it. A natural-sounding voice may have no corresponding act of speaking. A fluent paragraph may not have been composed sentence by sentence by the person or organisation presenting it. A video may combine recorded and generated elements in ways that are not immediately visible to its audience.

The familiar question, “What am I looking at?” is consequently accompanied by another: “How did this come to exist?”

The second question is increasingly difficult to answer merely by examining the first.

More than a distinction between real and fake

One consequence of this change is that familiar labels can become inadequate. AI-generated content is sometimes discussed through a simple opposition between “real” and “fake”. That distinction may be relevant in some circumstances, particularly where content purports to document something that did or did not happen. But it does not describe the full environment.

An AI-generated image might openly illustrate an imaginary city. Nothing deceptive necessarily follows from the fact that the city never existed. A synthetic voice might be deliberately used for narration and clearly presented as such. AI-generated text might accurately describe factual information even though no human individually composed every sentence. Conversely, a completely human-written statement can be false, and an unmodified photograph can be presented with a misleading description.

The method by which content was produced and the truth of what that content communicates are therefore related in some circumstances but are not the same question.

This boundary matters. The detailed examination of factual truth and verification belongs elsewhere within the AISF AI Guidance series. The concern here is more foundational: what happens to ordinary interpretation when the nature and origin of content can no longer be readily inferred from what people perceive?

This question applies even when nobody is attempting to deceive anyone.

Creation is becoming less visible in the finished product

Traditional content often carried clues about its production. A hand-drawn illustration looked different from a photograph. Early computer-generated imagery was often recognisable as computer-generated. Synthetic speech could sound mechanical. Automated text frequently displayed limitations that distinguished it from ordinary human writing.

Such distinctions have become less dependable as generation technologies have developed.

This does not imply that AI-generated content is always indistinguishable from human-created material. Quality varies greatly. Outputs may contain inconsistencies, errors, unusual features or obvious signs of artificial production. Different systems and forms of content have different capabilities and limitations, and these continue to change.

The broader development is that perceptual recognition alone cannot always be assumed to reveal origin.

That change affects more than deliberately deceptive content. It also applies to routine editing and assistance. A photograph might originate with a camera but later be expanded, reconstructed or altered by AI. A human writer might develop an argument while using AI to restructure passages or generate portions of text. A video might combine recorded footage with generated backgrounds or modified speech. Music might involve human performance alongside synthetic components.

In such circumstances, asking whether something is simply “human-created” or “AI-generated” may conceal more than it reveals.

Content increasingly has a production history rather than a single obvious category.

Later chapters will examine this continuum in greater depth. At this stage, its significance lies in recognising that the weakening relationship between appearance and origin is not confined to wholly synthetic creations. It also emerges from the many ways in which human and machine contributions can coexist within the same piece of material.

Different content forms, a common change

The implications also differ according to what people encounter.

Text occupies a distinctive position because readers have rarely been able to observe its creation directly. A written sentence has never physically demonstrated who composed it. Authorship has usually been understood through attribution and context rather than appearance alone. AI extends this existing separation by making fluent machine-generated language increasingly commonplace.

Images have a different history. Photography has often carried an intuitive association with physical observation: something was in front of a camera. That association has already been weakened by decades of digital manipulation, but generated imagery introduces the possibility that no corresponding photographed scene existed at all.

Audio raises another set of assumptions because voices are closely associated with identity. People routinely recognise relatives, colleagues, public figures and others through how they sound. Synthetic and modified voices complicate the relationship between hearing a recognisable voice and knowing that the person concerned actually spoke the words.

Video combines several of these expectations. Movement, voice, facial expression and environmental detail can together create a strong impression of witnessed reality, even though video has always been subject to staging, editing and manipulation.

These differences are important and will require separate examination. Yet across all four forms—text, image, audio and video—the underlying environmental change is similar: **the content people perceive may disclose progressively less about the process that produced it.**

An environmental change rather than an isolated technology

It would be possible to understand AI-generated content simply as another category of digital technology. From the perspective of everyday human experience, however, its significance may lie less in any individual tool than in its growing presence throughout the information environment.

People may encounter generated or modified material without deliberately seeking it. It can appear in search results, social feeds, advertisements, entertainment, workplace documents, educational materials, messaging platforms and public communication. Some material may be clearly labelled. Some may be recognisable from context. Some may involve only limited AI assistance. Other material may provide little indication of its production history.

This makes AI-generated content an environmental issue in the informational sense. The change is not restricted to people who actively use generative AI. It can affect anyone who reads, watches or listens within spaces where generated material circulates.

Nor will every encounter carry the same significance. Whether a synthetic image matters depends partly on what the image is doing. An openly fictional illustration, an advertisement, a historical reconstruction and an image presented as evidence of a current event occupy very different contexts. The technology involved may be similar while the interpretive consequences differ substantially.

Context therefore remains essential.

From automatic assumption to unresolved origin

For ordinary people, perhaps the most consequential change at this early stage is not that every piece of content must now be regarded with suspicion. Such a conclusion would replace one unreliable assumption with another.

Rather, origin is becoming a question that may sometimes remain unresolved.

A person can encounter content that appears entirely ordinary while lacking enough information to determine whether it was recorded, written, generated, modified or produced through some combination of those processes. In many everyday circumstances, that uncertainty may have little practical significance. In others, origin may substantially affect how the content is understood.

The significance depends upon context, purpose and consequence.

This is why the changing content environment cannot be reduced to a technical contest between increasingly capable generation systems and increasingly capable detection systems. Detection has an important place, but the broader human issue begins earlier. It concerns the assumptions people bring to information before any formal verification process occurs.

For generations accustomed to photographs, recordings, written authorship and moving images, those assumptions were shaped by technologies whose production processes were often imperfectly understood but broadly familiar. Generative AI introduces another layer between what is perceived and how it was produced.

The result is not the disappearance of reality, nor the end of meaningful communication. It is a change in one of the background conditions under which communication is interpreted.

Seeing something may still tell us a great deal about what it depicts.

Hearing something may still carry meaning, emotion and information.

Reading something may still communicate knowledge, ideas and experience.

But increasingly, none of those acts necessarily tells us, by itself, **how the content came into existence**.

That is the starting point for understanding what it means to live with AI-generated content.

Chapter 2 — From Human-Created to Machine-Generated: A Growing Continuum

The distinction between human-created and AI-generated content can appear straightforward. Something was either made by a person or generated by artificial intelligence. For some content, that description may be broadly accurate. Yet an increasing amount of material does not fit comfortably into either category.

A person may write a document and use AI to improve its wording. A photographer may take an original photograph and use AI-assisted tools to alter its background. A designer may generate several images and then substantially edit one of them. A speaker may record their own voice while software modifies its sound. A video may combine recorded footage, conventional editing and generated elements. A writer may develop an argument, ask an AI system to suggest alternative formulations, reject most of them and incorporate a few into the final work.

In each case, asking whether the finished content is “human-created” or “AI-generated” may produce an answer, but not necessarily a very informative one.

The changing content environment is therefore better understood not only through categories but also through a **continuum of human and machine involvement**. At different points along that continuum, AI may assist, modify, transform or generate. Human involvement may range from detailed creation and editing to broad direction, selection or approval. The resulting content may contain contributions from both that cannot easily be separated by looking at the final product.

Recognising this continuum does not require abandoning useful distinctions. It requires understanding what those distinctions can and cannot tell us.

Content creation was already a layered process

Even before contemporary generative AI, the idea of a single creator could simplify how content was actually produced.

A published article might involve a writer, editor, copy editor, photographer and designer. A film could involve hundreds of people as well as extensive digital processing. Advertising might combine photography, illustration, editing, typography and computer-generated imagery. Recorded music could pass through multiple stages of performance, mixing, correction and production. Software already automated spelling checks, image enhancement, audio processing and many other tasks.

Content creation has therefore long involved tools, collaboration and transformation.

What generative AI changes is not the existence of technological assistance but the range of contributions that technology can make. Earlier tools often operated primarily on material

supplied by people: correcting, arranging, processing or enhancing it according to relatively defined functions. Generative systems can also produce new text, imagery, sound and other material in response to human instructions or contextual inputs.

The boundary between tool and contributor consequently becomes less obvious.

This does not establish that AI should be regarded as an author, artist or creator in the same human sense. Those terms carry cultural, philosophical and legal meanings that vary across societies and contexts. The more immediate point is that the production process itself can now contain forms of machine generation that resemble activities previously associated mainly with human content creation.

Assistance can take many forms

At one end of the continuum, AI involvement may be limited.

Someone might write an entire passage and use an AI-enabled spelling or grammar function to identify possible errors. A photograph might receive automated adjustments to brightness or sharpness. Background noise might be reduced in an audio recording. A video platform might automatically generate captions.

Many people may regard the underlying material in such circumstances as substantially human-created, even though automated systems contributed to the finished form.

Further along the continuum, AI can become more involved in shaping content. A writer may ask for alternative ways to structure a paragraph. An image may be extended beyond its original frame using generated material. A voice recording may be translated while retaining characteristics of the original speaker. Portions of a video may be reconstructed or altered.

Further still, a person may request an entire initial output from a generative system and then edit it extensively. Alternatively, the generated output may be accepted with little modification.

Even these examples do not produce a universal scale with precise boundaries. The significance of each form of involvement depends upon the content, the technology, the purpose and the context in which the material appears.

A minor automated change may sometimes matter greatly. A substantial generated contribution may sometimes matter very little. The percentage of content produced by AI is therefore not, by itself, a complete measure of its significance.

Human direction does not always reveal human contribution

Generative AI also complicates the idea of human involvement because people can participate at very different levels.

One person might spend hours developing instructions, supplying reference material, evaluating outputs, revising details and assembling a final work. Another might enter a short

request and accept the first result. Both may technically have initiated AI generation, yet the nature and extent of their contribution are very different.

Similarly, human selection can matter even where the underlying material is generated. Choosing one image from dozens, deciding which passage to retain, arranging generated elements or rejecting unsuitable outputs all involve human judgement. But selection is not necessarily equivalent to producing the selected material.

This creates a distinction between **initiating, directing, generating, selecting, modifying and presenting content**.

Those activities can be distributed differently from one situation to another. The person who asks for content may not be the person who edits it. The organisation that publishes it may not know every detail of the tools involved in producing it. A platform may apply automated modifications after material has been uploaded. Several people and systems may contribute at different stages.

The finished content can therefore conceal a complex production history.

Chapter 1 established that appearance increasingly reveals less about origin. The continuum adds another dimension: even when AI involvement is known, that fact alone may reveal little about **how much, where or in what way** AI contributed.

The continuum differs across text, images, audio and video

The balance between human and machine involvement does not appear identically across content forms.

In text, AI might correct grammar, rewrite individual sentences, summarise human-written material, suggest an outline, generate sections or produce an entire draft. Human revision might then range from minor editing to substantial reconstruction. The final text does not ordinarily contain a visible record of which words came from which stage.

Images can involve similarly layered processes. A photograph may begin as a conventional camera capture but later include generated additions. A wholly generated image may subsequently be retouched manually. Several photographs and synthetic elements may be combined. What appears to be a single image may therefore contain elements with different origins.

Audio may combine recorded human speech with automated cleaning, generated words, synthetic translation or voice transformation. Music may combine human performance with generated or digitally modified components.

Video can bring many of these layers together. Recorded footage, generated imagery, synthetic or modified voices, conventional visual effects and human editing can coexist within the same production.

These variations make a single universal label increasingly difficult to apply with precision.

The issue is not that categories such as “AI-generated” have become meaningless. They can remain useful descriptions. The issue is that their meaning depends on what they are being used to describe.

Labels can simplify complex production histories

As generated content becomes more common, labels may be used to distinguish it from other material. Such labels can provide information about origin or production. But the continuum creates a question about what any label encompasses.

Does “AI-generated” refer only to material substantially produced by a generative system? Does it include AI-assisted editing? What about an original photograph whose background has been generated? Is the entire image AI-generated, partly generated or AI-modified? Does an automatically translated recording become synthetic audio if the speaker's voice is computationally reproduced in another language?

There is no need for this Guide to establish universal answers to those questions. Terminology, technical practices, platform policies and regulatory approaches can differ and may continue to develop.

What matters for understanding the environment is that **a label compresses a production history into a category.**

Sometimes that compression may be sufficient for the purpose at hand. In other situations, the distinction between generated, modified and assisted content may be important. Context determines how much information about production is relevant.

This also explains why absence of an “AI-generated” label cannot automatically establish that no AI was involved. If AI operates across a continuum, the threshold at which its involvement becomes significant enough to describe or disclose is itself a contextual question.

AI involvement and content quality are separate questions

The continuum can also be obscured when assumptions about origin become assumptions about quality.

Human-created material can be careful or careless, original or derivative, accurate or inaccurate. AI-generated material can likewise vary substantially in usefulness, coherence and relationship to underlying information. Mixed human-AI content adds further variation.

Knowing that AI contributed to something therefore does not, by itself, determine whether the finished material is good, bad, reliable, unreliable, valuable or trivial.

Nor does extensive human involvement automatically settle those questions.

This distinction is particularly important because discussions of AI-generated content can easily move from describing how something was produced to judging whether it should be trusted. Trust is an important subject, but it belongs principally to another part of the AIGuide

series. Detailed questions about whether factual claims are true similarly belong to the Guide concerned with truth and verification.

The focus here remains on understanding the nature of the content environment.

From that perspective, the continuum matters because it weakens a simple classification system that might otherwise appear sufficient for navigating that environment.

Different contexts may draw the boundary differently

The significance attached to human and AI involvement can also vary across societies, professions and forms of communication.

In some creative contexts, the process by which something was produced may itself be important to how audiences understand the work. In routine administrative communication, automated assistance may attract comparatively little attention. In journalism, education, advertising, entertainment, research or personal communication, expectations surrounding authorship and production may differ.

Those expectations are not necessarily fixed or globally uniform.

Cultural understandings of authorship, originality, creativity and acceptable technological assistance can vary. Institutions may develop different conventions. Professional fields may distinguish among forms of AI involvement according to their own purposes. Legal and regulatory classifications may also differ between jurisdictions.

A globally applicable understanding therefore needs to avoid assuming that one threshold between “human” and “AI” will have the same significance everywhere.

The technological continuum may be broadly recognisable while its social interpretation remains contextual.

Beyond the binary

The familiar human-versus-machine distinction remains intuitively powerful because it offers clarity. Yet the production of contemporary content can involve both human agency and machine generation in changing combinations.

A human may originate an idea while AI generates its initial expression. AI may propose material while a human substantially transforms it. A recording may preserve an original human performance while altering particular elements synthetically. A generated image may be incorporated into a larger work produced through human direction and conventional tools.

None of this requires the distinction between human and machine contributions to disappear. On the contrary, understanding those contributions may become more important precisely because they are increasingly intertwined.

The central change is that content does not always occupy one of two clean categories.

Between material created without generative AI and material generated substantially by AI lies a widening range of assistance, modification, transformation, generation, selection and human revision. The position of any particular item within that range may be clear, partially known or impossible for an ordinary observer to determine.

That uncertainty is part of the emerging information environment.

Chapter 1 began with the recognition that content increasingly does not reveal its origins merely through appearance. The continuum developed here shows why the question of origin itself is becoming more complex. It may no longer be enough to ask whether something came from a human or from AI.

Increasingly, the more revealing question is:

What roles did humans and machines each play in bringing this content into existence?

Understanding that question provides the foundation for examining how the continuum presents itself differently across text, images, voices and video.

Chapter 3 — Text, Images, Voices and Video: Different Forms, Different Signals

AI-generated content is often discussed as though it were a single type of material. In everyday experience, however, reading a passage of generated text is not the same as seeing a synthetic image, hearing a familiar voice or watching apparently realistic video. Each form of content carries different expectations about how it came into existence and different relationships with human perception.

These differences matter because people do not interpret all media in the same way. Written language has traditionally depended heavily on attribution and context. Photographs have often carried an intuitive association with something that existed before a camera. Voices can be closely connected with personal identity. Video combines image, movement and often sound in a form that can create a particularly strong impression of witnessed events.

Generative AI can operate across all of these forms. Yet its significance cannot be understood simply by saying that each can now be artificially produced. The more important question is what assumptions people have historically attached to each form and what happens as those assumptions become less dependable.

Chapter 1 established that content increasingly may not reveal its origins. Chapter 2 showed that human and machine contributions can exist along a continuum rather than within two clean categories. This chapter adds a further layer: **the signals people have traditionally used to interpret origin differ according to whether the content is text, image, audio or video.**

Text has always concealed much of its production

Written language provides an important starting point because text has never visually demonstrated its authorship.

A sentence printed on a page does not look different because one person rather than another wrote it. Readers ordinarily infer authorship from surrounding information: a name on a book cover, a newspaper byline, an email address, an institutional website or the circumstances in which a message was received. Anonymous and pseudonymous writing have existed for centuries, as have ghostwriting, collaborative authorship, editing, copying and false attribution.

The arrival of generated text therefore does not remove a visible connection between author and words in quite the same way that synthetic imagery can remove the connection between photograph-like content and a photographed scene. That visible connection was never strong in text.

What changes is the scale and ease with which coherent language can be produced without someone composing each sentence in the conventional manner.

A person can request a description, explanation, letter, summary or narrative and receive fluent text within moments. The resulting words may subsequently be edited extensively, lightly

modified or used substantially as generated. As Chapter 2 examined, these variations create different combinations of human and machine contribution.

For the reader, however, those production differences may leave no obvious visual trace.

Certain patterns of language may sometimes appear characteristic of machine generation, just as particular writing styles may appear characteristic of particular people. But style alone is not a reliable account of production history. Human writing varies enormously, generated writing varies across systems and instructions, and human editing can alter generated material.

The interpretive challenge surrounding text is therefore closely connected to **authorship and process**. Who is represented as speaking? What role did a person have in forming the words? Does the context imply human composition when the production process was more complex?

These questions are distinct from whether the statements themselves are factually correct. Human-written text can contain falsehoods, and generated text can contain accurate information. The Guide's concern at this point remains origin and interpretation rather than detailed factual verification.

Images carry a different historical expectation

Images, particularly photographs, have a different relationship with perceived reality.

People have long known that photographs can be staged, selected, cropped, retouched and manipulated. Photography has never provided an infallible record of events. Nevertheless, an ordinary photograph has traditionally implied that light from some physical scene reached a camera or comparable recording device. Something existed before the lens, even if the resulting image presented it selectively or misleadingly.

Synthetic imagery changes that underlying assumption.

An image can now resemble a photograph without corresponding to any photographed scene. A person depicted in it may never have existed. A landscape may combine plausible features into a location found nowhere in the physical world. An apparently ordinary object may have been generated rather than photographed.

At the same time, the continuum described in Chapter 2 complicates the distinction. A genuine photograph may contain generated additions. Objects may be removed or replaced. Backgrounds may be expanded. Lighting or details may be reconstructed. A generated image may itself later undergo conventional editing.

The finished picture therefore may not reveal whether it originated from a camera, from generation, or from a combination of both.

This does not make every synthetic image deceptive. An illustration for a fictional story, a conceptual design or an openly imaginative scene may never ask viewers to believe that it documents a real event. The interpretive significance changes when photograph-like appearance carries an implied connection with physical reality.

The important shift is therefore not simply that artificial images exist. Artificial imagery has existed in many forms for a long time. It is that **photographic appearance and photographic origin are increasingly separable**.

Voices connect content with identity

Audio introduces another set of expectations.

Human beings are accustomed to recognising one another through voice. Tone, rhythm, accent and other characteristics can become deeply familiar. A person answering a telephone call from a family member may recognise the caller before hearing a name. Public figures may likewise have widely recognisable voices.

Recorded audio has traditionally preserved a strong conceptual link between a voice and an act of speaking, even though recordings could be edited or manipulated.

Synthetic voice technology complicates that relationship.

Speech can be generated with characteristics resembling those of a human voice. Existing recordings can be modified. Translation may allow words to be presented in another language while retaining aspects of the original speaker's vocal identity. Generated narration can sound increasingly natural without corresponding to a person speaking those words at that moment.

Here again, the categories are not necessarily simple. A recording may begin with genuine human speech and then undergo automated cleaning. Portions may be altered. A person's real voice may form the basis for synthetic production. Human and generated speech may exist within the same piece of audio.

The significance of these possibilities is partly psychological and social. Voice is not merely a medium for carrying words; it can also function as a signal of identity, presence and familiarity.

Hearing a recognisable voice has therefore often supported an immediate assumption: **this person is speaking**.

As synthetic audio becomes more capable, that assumption may no longer follow automatically from sound alone.

This does not mean that familiar voices cease to have meaning or that every recording becomes uncertain. It means that vocal resemblance and the act of speaking can increasingly be separated technologically.

Video combines several signals at once

Video can present an even more layered interpretive environment because it brings together multiple forms of information.

A person may be visible, moving, speaking and interacting with an apparent physical setting. Facial expression, body movement, voice, background and sequence can reinforce one another.

The result often carries a strong sense that the viewer is observing something that occurred before a camera.

Yet video has never been a transparent window onto reality. Scenes can be staged. Footage can be selectively edited. Events can be rearranged. Visual effects have long created convincing fictional environments and actions.

Generative AI expands the ways in which these elements can be produced or modified.

Faces may be altered. Speech and mouth movements may be changed. Backgrounds may be generated. Existing footage may be transformed. Entire sequences may be created synthetically. Recorded and generated components can coexist within one piece of content.

As with images and audio, this does not mean that synthetic video is inherently misleading. Film, entertainment, education, simulation and creative expression may use generated video without representing it as documentary evidence.

The interpretive issue emerges from the gap between **what video appears to record and how its constituent elements were actually produced.**

Because video can combine visual and auditory signals, apparent consistency among those signals may itself become less informative about origin. A matching voice, face and movement may once have seemed mutually reinforcing evidence of a person's presence. Increasingly, technological processes can generate or modify several of those elements together.

Different forms can also converge

The distinction among text, image, audio and video remains useful, but digital content increasingly crosses those boundaries.

A text instruction can produce an image. Written material can become synthetic speech. A still image can contribute to generated video. Recorded speech can be transcribed, translated and reproduced in another language. Several forms can be assembled into a single piece of communication.

A social-media post, for example, might contain generated text, an altered image and synthetic audio within a short video. The audience encounters one item, but its production history may involve several different generative processes.

This convergence reinforces the limitations of judging origin from a single surface characteristic.

It also demonstrates why the changing information environment cannot be understood only by learning to recognise one particular form of generated material. The technologies, outputs and conventions surrounding them continue to develop. A visual characteristic associated with generated imagery at one point may become less common later. Particular speech patterns may change as systems improve. Linguistic tendencies may differ between systems and uses.

A durable understanding therefore needs to focus less on temporary markers and more on the changing relationship between **perception and production**.

Signals are not the same as proof

People naturally use signals to interpret what they encounter. An unusual visual detail, an unnatural voice, unexpected wording or inconsistent movement may cause someone to wonder how content was produced. Context can provide further clues. Labels, source information and surrounding circumstances may also contribute to interpretation.

But noticing a signal and establishing origin are different activities.

This distinction is important because AIGuide 10 is not intended to become a technical detection manual. The objective is not to provide a catalogue of supposed signs through which readers can classify content as human or synthetic. Such an approach would also risk becoming rapidly outdated as technologies change.

Instead, the broader observation is that the signals historically associated with different media are changing in informational value.

A photograph-like appearance does not necessarily establish photographic origin. A recognisable voice does not necessarily establish that the recognised person spoke the words. Natural movement in video does not by itself establish that the depicted event occurred before a camera. Fluent prose does not disclose how much human composition lies behind it.

At the same time, none of those observations establishes that the content is synthetic.

Uncertainty operates in both directions.

What remains is context

When surface characteristics disclose less about production, other forms of information can become more significant to interpretation. Where content appears, how it is presented, what it claims to represent, whether its production is explained and what relationship exists between its source and audience can all influence how its meaning is understood.

The importance of such context varies greatly.

For an obviously fictional image, the distinction between photography and generation may have little bearing on the audience's understanding. For an audio recording presented as the words of a particular person, origin may be central to what the recording means. A generated passage used openly as part of an automated service occupies a different context from the same passage presented as someone's personal testimony.

The technology alone therefore does not determine the significance of synthetic content.

Text, images, voices and video each bring their own histories, conventions and perceptual expectations into the generative-AI environment. Those differences explain why the same

underlying technological development can carry different consequences across different forms of communication.

Yet they also converge on the central change established at the beginning of this Guide.

People continue to read words, look at images, recognise voices and watch moving pictures much as they have before. What is changing is the amount of information those experiences necessarily provide about the processes behind them.

The surface of content still communicates.

Increasingly, however, **the surface does not tell the whole story of its origin.**

Understanding that distinction prepares the way for a deeper question. If appearance and origin can be separated, and if human and machine contributions can be intertwined, then familiar ideas about whether something is “authentic” become more complicated.

That is the conceptual territory to which the Guide turns next.

Chapter 4 — What Does “Authentic” Mean Now?

The word *authentic* is often used as though its meaning were self-evident. A photograph may be described as authentic because it records an actual scene. A document may be considered authentic because it genuinely came from the person or institution named as its source. A work of art may be valued as authentic because of who created it. A personal message may feel authentic because it appears to express the sender's own thoughts and emotions.

These meanings overlap, but they are not identical.

The growing presence of AI-generated and AI-modified content makes those differences increasingly important. If an image accurately represents a real place but was generated rather than photographed, in what sense is it authentic or inauthentic? If a person develops an idea but uses AI to express it more clearly, does the resulting text authentically represent that person's thinking? If someone's voice is synthetically reproduced saying words that the person approved but never physically spoke, what exactly is authentic: the message, the voice, neither, or both in different senses?

There is no single answer that applies to every circumstance.

The significance of authenticity depends on what people believe a piece of content represents and why that relationship matters in the context in which it appears.

Understanding AI-generated content therefore requires separating several questions that ordinary language often combines.

Authenticity is not the same as truth

One of the most important distinctions is between authenticity and factual truth.

A genuine photograph can be authentic as a photograph while being used to support a false account of what happened. It may show a real event but be presented as having occurred at another place or time. A human may genuinely write a statement that contains factual errors. A real recording may preserve words that are misleading or untrue.

Conversely, synthetic material can communicate something factually accurate. A generated illustration may accurately represent the structure of an object. AI-generated narration may correctly describe an historical event. A synthetic reconstruction may be based carefully on established information even though the depicted scene was not literally recorded.

Origin does not by itself determine truth.

This distinction is particularly important within the architecture of the AISF AI Guidance series. AIGuide 10 examines what generated and modified content changes about the information environment, including how people understand authenticity and origin. Detailed factual verification belongs primarily to AIGuide 9.

The distinction allows a more precise question to be asked here. Rather than treating “authentic” as another word for “true”, it becomes possible to ask: **authentic in relation to what?**

The answer may concern source, authorship, experience, representation or production.

Authenticity of source

In some circumstances, authenticity concerns whether content genuinely came from the source to which it is attributed.

A message may be authentic because the identified person actually sent it. A statement may be authentic because it was genuinely issued by the organisation whose name appears on it. A recording may be authentic in the sense that it originated from the stated source.

AI complicates this relationship because the source of a communication and the producer of every element within it need not be the same.

An organisation might issue material containing AI-generated text while still genuinely standing behind the communication. A person might send a message drafted partly with AI assistance. A company might use a synthetic voice in an authorised communication. In each case, machine-generated elements can exist within content whose institutional or personal source is genuine.

The reverse is also possible. Content can imitate the appearance, language, image or voice of a genuine source without actually originating from that source.

These situations demonstrate that **source authenticity and production method are separate dimensions**.

Knowing that AI contributed to content does not establish that the claimed source is false. Knowing that content appears to come from a familiar source does not establish how it was produced.

Authenticity of authorship

Authorship introduces another layer.

Traditional assumptions about written content often associate authorship with composing words. Yet, as Chapter 2 examined, AI can participate in writing at many levels. A person may originate the ideas and write the text while using limited AI assistance. Another may provide detailed arguments and ask AI to formulate them. Someone else may request an entire draft and make only minor changes.

All three situations involve a human, an AI system and a finished text, but the nature of the human contribution differs substantially.

The question “Who wrote this?” may therefore conceal several questions: Who developed the ideas? Who composed the words? Who selected the material? Who revised it? Who takes responsibility for presenting it?

Similar questions arise outside writing. A designer can direct the generation of imagery. A filmmaker can incorporate generated scenes. A musician can combine human performance with synthetic elements. The person publicly associated with the final work may have exercised substantial creative judgement even when parts of the material were generated.

The purpose here is not to determine where authorship legally or philosophically resides. Different legal systems, cultural traditions and creative communities may approach that question differently.

The informational point is narrower: **the name associated with content may no longer describe its production process as completely as audiences once assumed.**

Authenticity of representation

Another meaning concerns whether content authentically represents what it appears to depict.

Consider an image of a real city. It might be a conventional photograph. It might be a photograph that has been substantially modified. It might be generated entirely from descriptions and reference material. Each could resemble the same place.

The question of authenticity depends partly on what the image claims to be.

If it is presented as an artistic interpretation, generation may be entirely consistent with its purpose. If it is presented as a photograph documenting what a particular street looked like at a particular moment, the relationship between the image and an actual recorded scene becomes much more significant.

Audio and video present similar distinctions. A synthetic voice might deliver words authorised by the person whose voice it resembles. The message may represent the person's approved position even though the person did not physically speak the recording. Alternatively, a synthetic voice might attribute words to someone without their involvement.

The sound could be technologically similar in both cases. The relationship between the representation and the represented person would be very different.

Authenticity therefore cannot always be determined from the media form itself. It depends on the relationship among **content, origin, representation and context.**

Authenticity of human expression

Generative AI also raises a less visible question: when does content authentically represent a person's own expression?

This issue can arise even where nobody is being impersonated.

People routinely use tools to improve communication. Editors reshape writing. Translators express ideas in another language. Templates structure correspondence. Speechwriters prepare words for others to deliver. None of these practices began with AI.

Generative systems add another form of mediation.

A person who struggles to express an idea may use AI to help articulate something they genuinely think. Another person may use generated wording because it is convenient while giving little attention to whether it reflects their own voice. Between these examples lies a wide range of possible relationships between individual intention and final expression.

The presence of AI alone cannot determine whether the communication is personally authentic. Neither can the absence of AI guarantee authenticity. A person can write words manually that do not reflect their actual beliefs, just as machine-assisted wording may accurately convey an idea they sincerely hold.

This illustrates why authenticity is partly relational rather than purely technological.

The relevant question may not be “Was AI used?” but “What relationship does this content have to the person, experience or intention it appears to represent?”

Modification makes the boundary still less obvious

The continuum between human-created and machine-generated material further complicates authenticity because content may begin with an authentic record and subsequently change.

A photograph can contain both captured and generated elements. A recording can preserve a real voice while replacing individual words. Video can combine original footage with synthetic backgrounds or reconstructed features. Text can move repeatedly between human writing and AI-assisted revision.

At what point does modification change the character of the original?

There is no universal threshold.

Minor adjustments may leave the underlying representation substantially unchanged. Other modifications can alter the meaning of content while preserving much of its original appearance. The significance depends not merely on the amount of alteration but on what was altered and why that feature matters.

Changing the brightness of a photograph and adding a person who was never present are both modifications, but their relationship with the depicted event is fundamentally different.

A numerical measure of how much content was changed would therefore not necessarily answer the authenticity question.

Meaning matters.

Context shapes expectations

The same production process can carry different implications in different settings.

Audiences watching a fictional film ordinarily understand that what appears on screen may involve actors, visual effects, generated environments and other forms of construction. Authenticity in that setting may concern artistic intention or attribution rather than whether every visible event physically occurred.

Expectations may differ when content is presented as journalism, personal testimony, documentary material, educational evidence or official communication.

Advertising, entertainment, creative expression, workplace communication and social media each bring their own conventions. Those conventions also differ across countries, communities and generations and may evolve as synthetic content becomes more familiar.

This is why authenticity cannot be reduced to a fixed technical classification applicable everywhere.

What matters partly depends on what audiences are reasonably being invited to understand the content as representing.

Provenance and authenticity are related, but distinct

Knowing more about how content came into existence can help clarify questions of authenticity. Information about source, creation and modification can reveal relationships that appearance alone does not.

That information is often described through the idea of **provenance**.

Provenance does not automatically settle whether content is authentic in every meaningful sense. A complete account of production might reveal substantial AI involvement while leaving open questions about authorship, representation or intention. But it can make the production history less opaque.

This distinction will become important in the next chapter.

For now, authenticity can be understood not as a single property that content either possesses or lacks, but as a set of relationships. A piece of material can be authentic in one respect and problematic in another. Its source may be genuine while parts of its production are synthetic. Its representation may be fictional while its attribution is entirely transparent. Its words may be generated while accurately expressing the position of the person presenting them.

The growth of AI-generated content does not make authenticity meaningless.

It makes precision about authenticity more important.

When appearance alone reveals less about origin, the question is no longer adequately captured by asking whether something is simply “real” or “fake”. More useful distinctions emerge when

considering what the content claims to represent, where it came from, who stands behind it, how it was produced and what expectations surround its use.

The changing information environment therefore does not remove the concept of authenticity.

It reveals that authenticity was always more complex than it sometimes appeared.

And as synthetic and human-created material increasingly coexist, understanding **where content came from and what happened to it along the way** becomes a more significant part of understanding what that content is.

Chapter 5 — Provenance: Understanding Where Content Came From

When people encounter content, attention naturally goes first to what the content says, shows or sounds like. A reader considers the words. A viewer looks at the image. A listener hears the voice. Yet as AI-generated and AI-modified material becomes more common, another dimension can become increasingly relevant: **the history behind the content**.

Where did it originate? Was it recorded, written, generated or assembled? Has it been altered? Did several people or systems contribute to it? Is the version being encountered the same as the version originally produced?

Questions such as these concern **provenance**.

The term is used in different fields with different degrees of technical precision. In the context of this Guide, provenance can be understood broadly as information about the **origin and production history of content**. It concerns where material came from and, where relevant, what happened to it between creation and the point at which someone encounters it.

Provenance is not new. Museums have long been interested in the histories of artworks and objects. Researchers document sources and methods. Publishers maintain records of authorship and editions. Photographs, documents and digital files can carry information about when or how they were created. Institutions establish records showing where information originated and how it moved through processes.

What changes in an environment of increasingly synthetic content is the potential significance of provenance for ordinary interpretation.

When appearance alone provides less information about how something came into existence, information about its history can provide a different kind of context.

From the finished object to its history

A piece of content can be understood in at least two ways.

The first concerns what is immediately present: the words on the screen, the image being viewed, the sound being heard or the video being watched.

The second concerns the chain of events that produced that material.

For a traditional photograph, that chain might include a photographer, camera capture, editing and publication. For contemporary digital imagery, the history might instead involve an initial photograph, AI-assisted modification, generated additions, further human editing and eventual distribution. A generated image may begin with human instructions, pass through a generative system, undergo manual alteration and later be incorporated into another work.

Text can have an equally layered history. Someone may develop an idea, use AI to generate a draft, rewrite portions, send the material to another person for editing and publish the final version. Alternatively, the original text may be human-written but subsequently summarised, translated or reformulated through AI.

Audio and video can involve still more combinations.

The final item does not necessarily display this history.

This is one reason provenance becomes conceptually important. It shifts attention from asking only **“What does this look or sound like?”** to also asking **“What path brought it here?”**

The second question can reveal aspects of content that the first cannot.

Provenance is broader than identifying AI

It would be easy to treat provenance simply as a way of answering whether something was generated by AI. That would be too narrow.

The production history of content can include many elements unrelated to AI: who originated it, when it was created, whether it was edited, whether it was copied from another source, how it was published and whether the version being viewed has changed.

AI adds new possibilities within that history, but it does not replace the rest of it.

Consider a photograph that began with an actual camera capture and later received generated additions. Knowing only that “AI was used” says little about what role it played. Conversely, describing the image simply as a photograph does not disclose what happened after the camera recorded it.

A more complete production history might distinguish between the original capture and subsequent transformation.

The same applies to writing. A statement labelled “AI-assisted” could involve a spelling correction, structural rewriting or substantial generation. The label communicates something, but it compresses potentially different processes into the same description.

Provenance can therefore be understood as richer than a binary classification.

It concerns the sequence of contributions and transformations that connect an encountered item with its origins.

Provenance and authenticity

Chapter 4 examined why authenticity can have several meanings. Content may be authentic in relation to its source, authorship, representation or human expression. Provenance is related to these questions because information about production can help clarify those relationships.

But provenance and authenticity are not identical.

A detailed production history might show that an image was entirely generated. That fact does not, by itself, establish whether its use is appropriate or misleading. If the image is openly presented as a fictional illustration, its synthetic origin may be entirely consistent with what audiences are being invited to understand.

Similarly, provenance might establish that a photograph came directly from a camera. That does not prove that the photograph is being presented with an accurate description or that the scene itself was not staged.

Provenance tells us something about **history and origin**. Authenticity concerns the relationship between content and what it purports to represent. Factual truth raises another set of questions again.

Keeping these concepts separate prevents provenance from being asked to establish more than it can.

A chain can contain many participants

One of the consequences of the continuum examined in Chapter 2 is that content may have no single simple point of creation.

A person might provide the initial idea. An AI system might produce material from that direction. Another person might modify it. A different system might translate or enhance it. A platform might automatically process it during publication. Someone else might extract a portion and circulate it separately.

Which stage represents the origin?

The answer depends partly on what is being asked.

The intellectual origin of an idea may differ from the origin of the particular words used to express it. The origin of a photograph may be the camera capture even though parts of the finished image were generated later. A video clip encountered online may originate from a longer recording but have acquired a different meaning after editing and redistribution.

This makes provenance resemble a **chain** rather than a single label.

The chain can include creation, modification, assembly, publication and redistribution. Different stages may matter in different contexts.

This complexity is not unique to AI. Books have editions. Photographs have been reproduced and altered. Film passes through extensive production processes. News reports may draw on multiple sources. Digital material can be copied almost instantly.

Generative AI adds further forms of transformation to an already layered information environment.

Provenance can be complete, partial or absent

In practice, people rarely receive a complete history of everything they encounter.

Sometimes the source and production process are clear. Content may be explicitly described as generated, edited or reconstructed. A publication may explain how an illustration was created. An organisation may identify the use of synthetic narration. Context itself may make the nature of the material obvious.

In other situations, only part of the history is known.

A viewer may know who published an image but not who originally created it. A reader may know that AI was involved in producing a document but not whether its role was minor or extensive. A video may be traceable to an earlier version without revealing how that version was made.

Sometimes provenance may be largely absent.

Content can be copied, detached from its original context, reposted, compressed, reformatted or transmitted through channels that preserve little information about where it began. Screenshots can separate material from the page on which it appeared. Extracted clips can circulate without the surrounding video. Images can be reproduced without captions or attribution.

The informational environment therefore contains content with very different levels of provenance visibility.

The absence of provenance does not itself establish that content is synthetic, false or misleading. It means that one potential source of contextual information is unavailable.

Technical information and human understanding

Digital technologies can carry information about files, creation processes and modifications. New approaches may also seek to preserve or communicate information about the origins of digital content.

Those developments are relevant to the broader environment, but this Guide does not need to become a technical handbook for metadata, authentication systems or content credentials.

The human-centred issue is simpler.

Information about provenance is useful only to the extent that it can be understood in relation to the question being asked.

A highly detailed technical record may contain information that an ordinary viewer does not interpret. A simple explanation such as “generated illustration”, “AI-modified photograph” or “synthetic narration” may sometimes communicate the relevant fact more effectively. In another context, such a description may be too broad to explain a significant transformation.

The amount and type of provenance information that matters therefore varies with context.

There is no assumption that every person needs a complete production history for every piece of content encountered. Much everyday communication proceeds without such information and may continue to do so.

The significance of provenance increases when understanding origin materially affects the meaning attributed to the content.

Provenance can also be separated from content

Even where information about origin exists, another challenge remains: provenance information itself has to travel with the material.

Digital content moves easily.

An image published in one context may be downloaded and uploaded elsewhere. A video may be edited into a shorter clip. Text may be copied into another document. Audio may be extracted from a longer recording. Screenshots may reproduce material while removing surrounding information.

Each movement can change what an audience knows about the content's history.

This creates an important distinction between **content having a provenance** and **the audience having access to that provenance**.

Everything has some history of production. The informational question is whether that history is visible, preserved, communicated or recoverable when someone encounters the resulting material.

AI-generated content makes this distinction increasingly relevant because the finished material may itself provide fewer obvious clues about the process behind it.

Different contexts give provenance different weight

As with authenticity, the importance of provenance is contextual.

In entertainment, an audience may have little need to know the detailed production history of every generated background. In an art exhibition, the method of production may itself be part of how a work is understood. In journalism, information about whether an image records an event or illustrates it synthetically may carry greater significance. In personal communication, whether a familiar person actually recorded a voice message may be central to the meaning attributed to it.

Educational, commercial, institutional and creative settings may each develop different expectations about how production histories are described.

Those expectations can also vary across societies and change over time.

A globally applicable understanding of provenance therefore cannot depend upon one universal convention. The more durable concept is the relationship between **what people encounter and what they know about how it came to exist**.

Origin as part of meaning

For much content, production history will remain secondary. People read novels without knowing how every sentence was revised. They watch films without tracing the origin of every visual element. They encounter routine communications without reconstructing the tools used to prepare them.

AI does not necessarily change that.

What it changes is the range of circumstances in which origin may matter while remaining invisible from the finished content itself.

A photograph-like image may or may not have begun with a photograph. A familiar-sounding voice may or may not originate from an act of speaking. A fluent passage may contain varying degrees of human and machine composition. A video may combine recorded and generated elements in ways that are not apparent from viewing it.

Provenance provides a way of thinking about this gap without assuming in advance that the content is good or bad, authentic or deceptive, true or false.

It asks instead about history.

Where did the content come from? What processes shaped it? What transformations occurred before it reached the person encountering it?

These questions become increasingly relevant as the surface of content reveals less about its creation.

They also prepare the ground for a more specific issue. When synthetic capabilities can reproduce not merely generic content but recognisable people, voices, places and apparently real events, questions of origin become connected with representation in particularly powerful ways.

That is the territory to which the Guide turns next.

Chapter 6 — Synthetic People, Voices, Places and Events

AI-generated content becomes particularly significant when it does more than create words, images or sounds. It can also represent **people, voices, places and events** that audiences may associate with the physical world.

An image may depict a recognisable person in circumstances that were never photographed. A voice may resemble someone who did not speak the words being heard. A generated streetscape may look like a real location without corresponding precisely to anything that existed there. Video may present an apparently coherent event even though some or all of its elements were generated or substantially modified.

These possibilities are often discussed through the language of deepfakes. That term captures an important part of the contemporary environment, particularly where realistic synthetic media represents identifiable people. But the wider subject is broader. Synthetic representation can include fictional people, reconstructed places, simulated events, translated voices, educational visualisations, entertainment, advertising and many other forms of content.

The defining issue for this chapter is therefore not a particular technology or category of manipulation. It is the changing relationship between **representation and the reality that audiences may assume lies behind it**.

The earlier chapters established that appearance increasingly reveals less about origin, that human and machine contributions can exist along a continuum, and that authenticity and provenance are distinct but related concepts. Synthetic representations bring these developments together in a form that can be especially powerful because people are accustomed to treating faces, voices, places and recorded events as signals of presence.

When a person can be represented without being present

Images and recordings have long allowed people to be represented when they are not physically present before an audience. A portrait preserves someone's appearance. A recording preserves a voice. Film allows actions to be viewed long after they occurred.

Traditionally, however, many such representations began with some form of physical presence. A person stood before a camera, spoke into a microphone or performed an action that was recorded.

Generative AI can loosen that connection.

A photograph-like image of a person can be created without that person posing for the depicted scene. Their appearance can be inserted into or removed from existing material. Facial expressions, clothing, surroundings and apparent activities can be generated or changed.

The significance of this depends greatly on context.

A clearly fictional or satirical image of a public figure may be understood by its audience as imaginative representation. A synthetic historical reconstruction may deliberately show a person in a scene for which no recording exists. Entertainment may create fictional characters who resemble no real individual.

Different questions arise when generated content appears to document something about a real person.

The image itself may look equally convincing in each case. What changes is the relationship between what viewers see and what they are being invited to believe occurred.

This is why synthetic representation cannot be understood solely through visual realism. **Realism describes appearance; it does not establish the relationship between the representation and an underlying event.**

A familiar voice without an act of speaking

Voice carries its own distinctive relationship with identity.

People can recognise one another through patterns of speech that are difficult to describe but immediately familiar. Accent, tone, rhythm, pitch and habitual expression can collectively create a strong sense that a particular person is present.

Synthetic voice technologies can reproduce aspects of that familiarity without requiring the person to speak the particular words being heard.

Again, several very different situations can result.

A person may authorise a synthetic version of their voice to narrate material they approve. A speaker may record words in one language and use technology to produce another-language version that retains characteristics of the original voice. Entertainment may use synthetic voices for fictional characters. Accessibility applications may allow someone to communicate through a digitally reproduced voice.

In such circumstances, synthetic speech is not necessarily attempting to misrepresent who stands behind the communication.

But similar capabilities can also create audio that attributes words to someone who did not say or approve them.

From the listener's perspective, the surface signal—the recognisable voice—may be similar.

This creates an important distinction between **vocal identity and communicative origin**. A recording can sound like someone without corresponding to an act in which that person spoke those words.

The distinction does not imply that voice recognition becomes useless. It means that hearing a familiar voice may provide less information about the production history of a recording than listeners have historically assumed.

Places can be plausible without being recorded

Synthetic representation also extends beyond people.

Generated imagery can depict locations ranging from imaginary landscapes to convincing versions of real cities. A scene may contain familiar architecture, road signs, vegetation, weather conditions and other details that collectively make it appear geographically plausible.

Yet plausible is not the same as recorded.

A synthetic scene may represent a place approximately. It may combine characteristics from several locations. It may depict a real location under conditions that never occurred. It may show a hypothetical future development, an historical reconstruction or an entirely fictional environment.

These possibilities can be useful in many contexts. Architecture, entertainment, education, planning and creative work have long relied on representations of places that do not claim to be documentary photographs.

Generative AI expands the accessibility and realism of such representations.

The interpretive issue arises when the distinction between **illustrating a place** and **recording a place** matters to what audiences understand.

A generated image showing how a proposed building might look performs a different function from an image presented as evidence that the building already exists. The visual realism of the two images may be comparable, but their relationship to physical reality is different.

Once again, context carries information that appearance alone may not.

Events present a deeper representational question

Events introduce another level of complexity because they involve not merely objects or identities but actions unfolding in time.

People have long understood that dramatic scenes can be staged and that films can create fictional events. Visual effects can show impossible actions. Documentary footage can be edited. News footage can be selectively framed.

Generative technologies add the ability to create or transform apparently recorded events with increasing flexibility.

A crowd may be generated. A person's movements may be altered. Weather or surroundings may change. A short sequence may be created without a corresponding filmed event. Recorded and synthetic elements may be combined so that the final content has no single simple relationship with what physically occurred.

The phrase “Did this happen?” can therefore conceal several possibilities.

Perhaps nothing depicted occurred. Perhaps an underlying event occurred but the representation changed important details. Perhaps the scene is a reconstruction intended to illustrate something known from other sources. Perhaps genuine footage contains only limited synthetic modification.

These distinctions matter because an event is not necessarily either wholly recorded or wholly invented.

The continuum examined earlier applies here as well.

Reconstruction and simulation have legitimate contexts

Synthetic representation should not be understood only through misuse.

Human societies have long reconstructed events that could not be directly observed. Historical documentaries use actors and visualisations. Scientific models depict processes that cannot easily be filmed. Educational materials illustrate past environments. Architects visualise buildings before construction. Entertainment routinely creates worlds, characters and events that audiences understand to be fictional.

AI can contribute to these existing practices.

A generated reconstruction may help audiences visualise a historical setting. A simulated environment may illustrate a scientific concept. Synthetic speech may make material accessible in additional languages. Creative works may depict people and places that have never existed.

The fact that content does not correspond to a directly recorded event therefore does not determine whether its use is appropriate, useful or misleading.

What matters to interpretation is partly the **relationship between the content's actual production and the relationship audiences understand it to have with reality**.

An openly presented reconstruction and an apparently documentary recording can contain similar synthetic techniques while communicating fundamentally different claims about themselves.

This reinforces a theme running throughout the Guide: technology alone does not determine informational meaning.

Real people create particular sensitivities

When synthetic content represents identifiable people, additional human considerations arise.

A person's face and voice are closely connected with how others recognise them. Generated representations can therefore create situations in which someone's apparent presence becomes separable from their actual participation.

This can affect public figures and private individuals alike.

The consequences may differ considerably. A synthetic representation may be humorous, artistic, authorised, commercial, educational, harmful or simply confusing depending on its context and use. Laws and social expectations concerning identity, consent, reputation and intellectual property also differ across jurisdictions.

AIGuide 10 does not provide legal guidance on these matters, nor does it attempt to prescribe how particular representations should be governed.

Its concern is the informational change underneath them.

Historically, seeing a person's face in apparently photographic material or hearing their voice in a recording could support a strong everyday inference that the individual had participated in producing the underlying representation. That inference can no longer be treated as universally dependable.

The person depicted and the person participating have become technologically separable.

Synthetic representation can affect genuine material too

The consequences do not end with generated content itself.

Once people know that convincing representations can be synthesised, their interpretation of genuine photographs, recordings and videos may also change.

An authentic recording may be questioned because a synthetic equivalent appears technically possible. A genuine image may be dismissed as generated. Uncertainty about production can therefore move in both directions: synthetic content may sometimes appear recorded, while recorded content may sometimes be suspected of being synthetic.

This broader consequence will be examined more fully later in the Guide.

At this stage, it demonstrates why synthetic representation is an environmental change rather than merely a new category of media. Its existence can alter the assumptions surrounding other content as well.

Representation is not presence

Faces, voices, places and moving images remain powerful forms of communication. They can convey experiences that text alone cannot easily reproduce. Synthetic technologies do not remove that value.

What they change is one of the assumptions that can accompany representation.

Seeing a person is not necessarily evidence that the person stood before a camera in the circumstances depicted. Hearing a voice is not necessarily evidence that the person spoke the particular words. Seeing a plausible location does not necessarily mean a camera recorded that scene. Watching a coherent sequence does not necessarily establish that the event occurred in the form presented.

None of these observations determines whether the content is truthful, deceptive, useful or harmful.

They establish something more basic.

Representation and physical occurrence can increasingly be separated while the representation itself remains convincing.

That separation makes origin, context and production history more significant to understanding what content represents.

It also creates a larger question for the information environment. Synthetic people, voices, places and events can be produced not merely one at a time but increasingly quickly, repeatedly and in many variations.

The significance of generated content therefore lies not only in what can be created.

It also lies in **how much can be created, how rapidly, and how easily different versions can enter circulation.**

That shift from individual representation to scale is the territory to which the Guide turns next.

Chapter 7 — When Content Can Be Produced at Scale

The significance of AI-generated content does not arise only from what an individual system can produce. It also arises from **how much content can be produced, how quickly it can be created and how easily variations can be generated and distributed.**

A single synthetic image may resemble something that could previously have been produced through photography, illustration or digital editing. A generated passage may perform a function once fulfilled by human writing. Synthetic audio or video may reproduce effects that were already possible through specialised production techniques. Viewed one item at a time, some of these developments can appear primarily as changes in the tools available for content creation.

Scale changes the picture.

Generative systems can reduce the time and effort required to produce certain forms of material. They can create multiple versions from a common starting point, adapt content for different audiences or contexts, and move between text, image, audio and video. These capabilities do not mean that production is effortless, unlimited or consistently successful. Quality, cost, computing resources, human oversight and technical limitations continue to matter.

Nevertheless, the amount of content that can potentially be produced by a person, organisation or automated process is changing.

That development matters even when the content itself is neither false nor deceptive. The information environment is shaped not only by whether individual items are accurate or authentic, but also by their **volume, velocity, variation and visibility.**

Content production has always been constrained

Producing and distributing information has historically required resources.

Writing a newspaper article required someone's time. Producing an illustration required artistic or technical effort. Recording professional audio required speakers, equipment and production. Film and video could require substantial teams, facilities and budgets. Printing and broadcasting introduced further costs of distribution.

Digital technologies changed many of these constraints long before generative AI.

Word processing accelerated writing and revision. Digital cameras removed much of the marginal cost of taking additional photographs. Editing software made sophisticated production accessible to more people. Social platforms allowed individuals to distribute material to potentially large audiences without owning printing presses or broadcasting infrastructure.

The internet therefore dramatically increased the scale at which content could already be created, copied and circulated.

Generative AI represents another change within this longer development.

Its distinctive contribution is not simply cheaper distribution or easier reproduction of existing content. It can reduce some of the effort involved in **producing new variations of content itself**.

One starting instruction can generate several drafts. An image concept can produce multiple visual interpretations. Text can be reformulated for different lengths or audiences. Material can be translated or adapted. Synthetic voices can render written content as speech. Increasingly, one content form can contribute to the generation of another.

The constraint does not disappear. It moves.

Human time may shift from producing every element directly toward specifying, selecting, reviewing, editing and deciding what to use.

More content does not mean more information

The ability to produce more material creates an important distinction between **content volume and informational value**.

Ten explanations of the same idea do not necessarily provide ten times as much knowledge. Hundreds of images representing similar concepts may increase visual material without increasing understanding. Repeated reformulations of an existing claim can make that claim more visible while adding little new evidence.

This distinction existed before generative AI. Digital duplication and social sharing already allowed the same information to appear repeatedly.

Generation introduces a further possibility: repetition no longer needs to involve exact copies.

The same underlying idea can be expressed through different wording. Similar images can be generated with different compositions. Messages can be adapted to different tones. Multiple videos or audio recordings can communicate substantially the same proposition while appearing to be separate pieces of content.

The information environment can therefore become more varied at the surface while remaining repetitive underneath.

For an ordinary person encountering this material, that difference may not always be obvious. Ten independently produced accounts and ten generated variations derived from a common source can look like ten separate items.

Their informational significance may be very different.

Variation can be produced rapidly

Generative systems are particularly suited to variation.

A piece of text can be shortened, expanded, simplified, reformulated or translated. An image can be recreated with different backgrounds, perspectives or visual characteristics. Audio can be rendered through different voices or languages. Similar video sequences can be produced with altered details.

Variation has many ordinary uses.

Organisations routinely adapt communications for different audiences. Educational materials may be expressed at different levels of complexity. Creative professionals explore alternatives before selecting a final version. Translation allows information to cross linguistic boundaries. Accessibility can require the same underlying material to appear in different forms.

Generative technologies can support these processes.

At the same time, variation complicates the way audiences perceive scale. Material that is substantively connected may appear diverse because its surface characteristics differ.

This matters because people often use the apparent number of separate sources, voices or examples as part of their informal understanding of what is prominent in the surrounding environment.

Generated variation can alter that relationship.

The presence of many distinct-looking items does not necessarily establish that many independent human processes produced them.

Production and distribution are separate forms of scale

It is also useful to distinguish between producing content and making that content visible.

Generative AI can increase production capacity, but producing a million items does not mean that a million people will see them. Distribution remains shaped by platforms, networks, search systems, recommendation processes, publishers, organisations, communities and individual sharing.

An item can be produced instantly and remain unseen.

Conversely, one piece of content can reach millions of people without being generated at scale because digital distribution can reproduce it almost without limit.

The contemporary environment therefore combines two different forms of amplification: **scalable production and scalable distribution.**

When the two coincide, large quantities of material can potentially be created and circulated rapidly. But the relationship is not automatic. Visibility depends on additional systems and human behaviour.

This distinction prevents the effects of generative AI from being overstated. Generation increases the potential supply of content. It does not determine what audiences encounter.

Yet when supply becomes easier to expand, the mechanisms that determine visibility can become correspondingly important to the character of the information environment.

Scale can change what receives attention

Human attention remains limited even when content production expands.

A person cannot read every article, inspect every image or watch every video available. Increasing content volume therefore does not increase the amount of attention available to receive it.

This creates an asymmetry.

Content can become easier to produce while human capacity to interpret it remains constrained by time, attention and context.

The result is not necessarily that people become unable to understand information. Rather, selection becomes an unavoidable part of the environment. Individuals, institutions and digital systems already filter large quantities of material in order to make information manageable.

AI-generated content enters this environment as another potentially expanding source of material.

Its significance therefore cannot be measured simply by counting how much is produced. What matters is also which material becomes visible, how often it appears, where it appears and what relationship it has to other content competing for attention.

These questions concern the structure of the information environment rather than the quality of any single generated item.

Speed changes the timing of content

Scale is also temporal.

Generative systems can shorten the interval between an event, idea or instruction and the production of content relating to it. Material can potentially be created while circumstances are still developing and revised repeatedly as new information appears.

Speed can have practical value. Communications can be prepared more rapidly. Material can be translated sooner. Organisations can respond to changing circumstances with less production delay. Individuals can explore ideas quickly.

But faster production also means that greater quantities of material can appear before an informational situation has stabilised.

This is not exclusively an AI phenomenon. Social media and continuous news cycles already accelerated the movement of information. Generative AI can add another layer by accelerating the production of material that enters those channels.

Speed therefore affects not only how quickly people receive information but also **how mature that information may be when content begins accumulating around it.**

Again, this observation does not determine whether the material is accurate. Detailed factual verification belongs primarily to AIGuide 9. The point here is that production velocity changes the conditions under which people encounter content.

Scale can operate without deception

It is particularly important not to treat large-scale generation as synonymous with manipulation.

A business may generate routine product descriptions. An educational organisation may produce multiple versions of explanatory material. A public institution may translate information into several languages. A creative enterprise may generate large numbers of preliminary concepts before selecting a small number for publication.

Scale can therefore result from ordinary attempts to communicate, create or operate more efficiently.

The informational consequences nevertheless remain relevant.

If many organisations and individuals adopt similar capabilities, the cumulative amount of machine-assisted material in circulation may grow even though no single participant intends to overwhelm the environment.

Large environmental changes can emerge from many ordinary decisions.

This is one reason the subject cannot be understood only by examining obviously harmful uses. The overall content environment may change through widespread routine adoption as well as through deliberate attempts at influence.

Scale can make origin harder to perceive

Earlier chapters examined how an individual piece of content may conceal its production history. Scale can intensify that uncertainty.

When large quantities of material appear across different forms and channels, audiences may have limited opportunity to consider the provenance of each item. Generated variations may also become detached from their original production context as they are copied, shared or transformed.

The question is no longer merely whether the origin of one image, passage or recording is visible.

It becomes whether an environment containing large volumes of mixed human-created, AI-assisted, modified and generated material provides enough contextual information for people to understand what kinds of content they are encountering.

The answer will differ across settings.

Some environments may establish clear conventions for identifying generated material. Others may contain little information about production. Many will include a mixture.

Global practices are also unlikely to develop uniformly. Platforms, institutions, industries and jurisdictions may take different approaches, while expectations among audiences may evolve at different rates.

From individual content to environmental conditions

The first six chapters of this Guide largely examined content at the level of the individual item: where it came from, how human and machine contributions can combine, what different media forms imply, how authenticity can be understood and how people, voices, places and events can be represented synthetically.

Scale requires a wider perspective.

When production becomes faster and more accessible, the significance of generative AI is not confined to whether a particular image is synthetic or whether a particular paragraph was machine-generated.

It concerns the composition of the environment through which people move.

More content can exist. More variations can be created. Similar ideas can appear in many forms. Production can respond rapidly to events and audiences. Digital distribution can then place some portion of this expanded supply into widespread circulation.

Human attention, however, remains finite.

That tension makes scale an informational issue in its own right.

It does not follow that an environment containing more generated content necessarily contains less valuable information. Nor does greater volume automatically produce confusion. New capabilities can broaden expression, accessibility and participation as well as increase quantity.

The more durable observation is that **when the economics and effort of producing content change, the environment created by that content can change with them.**

This leads to another question.

If large quantities of material can be generated across text, imagery, audio and video, familiar boundaries between different purposes of content can also become less obvious. Something may inform, illustrate, persuade, entertain or simulate—and sometimes perform several of those functions at once.

Understanding that blurring is the next stage in examining what it means to live within an increasingly synthetic information environment.

Chapter 8 — The Blurring of Information, Expression and Simulation

Content does more than communicate facts. It can explain, persuade, entertain, imagine, illustrate, simulate, provoke, document and express. These purposes have always overlapped, but people have developed conventions that help them interpret what they encounter. A news photograph is generally approached differently from an illustration. A fictional film is understood differently from documentary footage. An advertisement carries different expectations from a personal message. A simulation is not ordinarily treated as a recording of an event that actually occurred.

Generative AI does not remove these distinctions, but it can make the boundaries between them less immediately visible.

A realistic image may document something that happened, illustrate something that might happen or depict something that never happened. A synthetic voice may deliver factual information, perform a fictional role or represent the authorised words of someone who never physically spoke them. Video may combine recorded material with generated elements. Written content may move between explanation, imitation, fictionalisation and representation without any obvious change in its visual form.

The challenge is therefore not simply distinguishing human-created content from machine-generated content. It is also understanding **what kind of communicative object a person is encountering and what relationship that object is intended to have with reality.**

This distinction becomes increasingly important as the same generative capabilities can be used across information, creative expression and simulation.

Content has always depended on conventions

People rarely interpret content without context.

A painting of an historical event is not normally assumed to be a direct visual record of the event. A film based on real people may combine documented facts with dramatisation. A diagram can simplify physical reality in order to explain it. A political cartoon may deliberately exaggerate. Advertising frequently creates idealised representations. Fiction can incorporate real places and events without claiming that its story actually occurred.

These conventions allow audiences to understand material that would be misleading if interpreted literally but meaningful when interpreted according to its purpose.

The same principle applies to digital content.

Computer-generated imagery has been used extensively in film, games, architecture, scientific visualisation and advertising. Virtual environments have represented places that do not exist. Digital editing has transformed photographs and recordings. Simulation has allowed people to

explore events that cannot conveniently, safely or economically be reproduced in the physical world.

Generative AI therefore enters an information environment that already contains many forms of constructed representation.

What changes is the accessibility, flexibility and realism with which such representations can be produced across multiple media.

As the technical differences between documentary-looking and synthetic-looking material become less obvious, audiences may depend more heavily on contextual differences to understand what they are seeing.

Information can be generated without becoming fictional

The term “generated content” can sometimes create an impression that the resulting material is necessarily imaginary.

That is not the case.

Generated text can explain established information. A synthetic voice can narrate factual material. An AI-generated diagram can represent a real concept. A generated visualisation can illustrate a process that exists even though no camera could directly photograph it.

The production method and informational purpose are separate.

This is similar to earlier distinctions between origin, authenticity and truth. Something can be generated while intending to communicate information about the real world. Whether it does so accurately is a further question, but its generated origin alone does not make its purpose fictional.

This distinction matters because synthetic content may increasingly appear inside ordinary informational settings.

A reader may encounter generated summaries, illustrations, translations or explanations without leaving the broader environment of information. The material may not present itself as creative expression even though generative processes contributed to its production.

The boundary is therefore not between “generated” and “informational”.

Generated content can be informational, expressive, simulated or some combination of these.

Expression can resemble documentation

Creative expression creates a different relationship with reality.

Artists, writers, filmmakers and other creators have always imagined people, places and events that never existed. The ability to represent the unreal is fundamental to many forms of culture.

Generative systems extend the tools through which such expression can occur.

An imaginary scene can be rendered with photographic characteristics. Fictional dialogue can be spoken by natural-sounding synthetic voices. A written description can become moving imagery. Historical or futuristic settings can be produced without conventional filming.

None of this inherently creates confusion.

The interpretive issue arises because **the visual or auditory style of documentation can increasingly be used for content whose purpose is imaginative or expressive.**

A fictional scene does not need to look fictional.

This distinction existed before AI. Cinema could convincingly depict imaginary events, and digital effects could create realistic environments. Generative AI broadens access to comparable forms of representation and can reduce the resources required to produce them.

As a result, documentary appearance may become a weaker signal of documentary purpose.

A viewer may need context to understand whether a realistic scene is intended to record, reconstruct, illustrate, entertain or imagine.

Simulation occupies the space between representation and occurrence

Simulation adds another layer because it often represents something that did not literally occur while maintaining a deliberate relationship with reality.

A simulation may explore how a building could look before it is constructed. An educational reconstruction may depict an historical environment for which no visual recording exists. A scientific visualisation may represent processes derived from data or models. Training material may create hypothetical situations so that people can consider circumstances without experiencing the real event.

Such content is neither simply documentary nor simply fictional.

Its value may lie precisely in showing something that was not directly observed.

Generative AI can make simulations easier to produce and more visually or audibly convincing. Text, images, voices and video can be combined to represent hypothetical situations in increasingly natural forms.

This creates an important distinction between **a representation being realistic and the represented event being real.**

A highly realistic simulation can be useful because it resembles the circumstances it seeks to illustrate. Yet that same realism can make its status less obvious if the context identifying it as a simulation is removed.

Provenance therefore intersects with communicative purpose. Knowing that content was generated may tell an audience something about its production, while knowing that it is a simulation tells them something different about how it should be understood.

One item can perform several functions

The boundaries become still less clear when a single piece of content combines purposes.

An educational video might contain factual narration, generated visualisations and hypothetical scenarios. A news report might use an illustration to explain an event for which no footage exists. Advertising can combine factual product information with imaginative imagery. Documentary production may use reconstruction alongside original recordings. Entertainment may incorporate real archival material into fictional storytelling.

These combinations are not necessarily problematic. They can be effective forms of communication.

But they demonstrate why classifying an entire piece of content under one label may conceal important differences within it.

A video can contain recorded and generated material. It can inform and simulate. It can document one aspect of reality while illustrating another. The audience may understand those transitions clearly—or may not notice them at all.

The continuum described earlier in this Guide therefore applies not only to **how content is produced**, but also to **what content is doing**.

Informational, expressive and simulated functions can coexist within the same material.

Context can disappear during circulation

The distinction among these functions becomes particularly important when content moves away from its original setting.

A simulated scene may be clearly identified within a full documentary but lose that explanation when extracted as a short clip. An artistic image may originally appear with a description identifying it as generated but later circulate without the description. A fictional sequence may be reposted separately from the work that established its fictional setting.

Digital circulation can separate content from the contextual signals that initially explained its purpose.

This phenomenon is not unique to AI-generated material. Quotations, photographs and video clips have long been removed from their original contexts.

Synthetic content adds another dimension because the extracted material may itself provide few clues about whether it documents, reconstructs or imagines what it depicts.

The same realistic video could therefore be understood very differently depending on the information accompanying it.

The content has not necessarily changed.

Its **interpretive environment** has.

Purpose matters as much as production

The earlier chapters established that asking whether AI was involved does not answer every important question about content.

Chapter 8 adds another reason why.

Two items may be produced using similar generative technologies but perform entirely different functions. One may be a fictional scene. Another may be an educational reconstruction. A third may be presented as a record of an actual event. A fourth may advertise a product through an idealised representation.

Knowing the production method helps explain origin.

Knowing the purpose helps explain meaning.

Neither necessarily establishes factual accuracy, authenticity in every sense or the appropriate level of trust. Those questions involve additional considerations and, where detailed verification is required, territory belonging primarily to other Guides in the AISF series.

For AIGuide 10, the important observation is that an increasingly synthetic information environment requires distinctions not only between human and machine production but also among **different relationships between content and reality**.

Different settings carry different expectations

The significance of these distinctions varies globally and across social contexts.

Journalism, entertainment, education, advertising, art, workplace communication and personal communication do not necessarily share the same conventions. Expectations can also differ among jurisdictions, institutions, professions, cultures and generations.

A generated reconstruction that is readily understood within one setting may be unfamiliar in another. Audiences accustomed to extensive digital effects may interpret highly realistic creative imagery differently from audiences with other media experiences.

These differences caution against assuming that everyone will draw the same boundary between information, expression and simulation.

Nor are the boundaries fixed. Conventions develop as technologies become familiar. New labels, practices and expectations may emerge. Some distinctions may become clearer while others become less important.

The durable issue is not the precise convention adopted at any particular time.

It is the continuing need to understand **what relationship content claims to have with the world it represents.**

When categories become less visible

The growth of AI-generated content does not mean that information, creative expression and simulation collapse into one indistinguishable category.

They remain conceptually different.

What can become less obvious is which category—or combination of categories—an individual item occupies when its appearance alone does not disclose its purpose or production.

A photograph-like image can be documentation, illustration or imagination.

A natural-sounding voice can preserve an act of speaking, deliver authorised synthetic narration or perform fiction.

A realistic video can record, reconstruct, simulate or invent.

A fluent passage can report, explain, imitate or imagine.

Understanding the content therefore increasingly involves more than recognising its surface form.

Origin matters. Provenance matters. Context matters. Purpose matters.

These factors do not replace human interpretation. They become part of what interpretation must sometimes accommodate.

The consequence is not necessarily permanent uncertainty. Many contexts remain clear, and audiences routinely understand constructed content without difficulty. The deeper change is that **realistic appearance alone is becoming less capable of signalling what kind of relationship content has with reality.**

That change becomes still more significant when content is produced at scale and repeatedly encountered across different settings.

Repetition can make material feel familiar. Apparent prevalence can shape perceptions of what is common or widely accepted. Large quantities of similar content can create impressions that are not explained merely by examining one item at a time.

The next stage of the Guide therefore turns from the purpose of content to another feature of the emerging environment: what happens when familiarity and repetition themselves become part of how credibility is perceived.

Chapter 9 — When Genuine Content Can Also Be Doubted

The most obvious consequence of convincing AI-generated content is that synthetic material may sometimes be mistaken for something recorded or human-created. Yet there is another consequence that operates in the opposite direction.

Genuine content can also become easier to doubt.

A real photograph can be described as generated. An authentic recording can be dismissed as synthetic. Genuine video can be questioned because people know that similar-looking material could, in principle, be created or altered using artificial intelligence. Human-written text may be suspected of machine generation because of its style or fluency.

This creates an important shift in the information environment. Synthetic content does not affect only the material that AI actually generates. Awareness of synthetic capability can influence how people interpret content that has not been generated at all.

The distinction matters because uncertainty can now move in two directions. People may mistake synthetic material for recorded or human-created content, but they may also mistake genuine material for synthetic content.

The result is not that everything becomes equally uncertain. Most everyday communication continues to occur within contexts that provide substantial information about origin and meaning. Rather, a new possibility enters ordinary interpretation: **realism is no longer necessarily sufficient to settle whether something is real, while suspicion is no longer necessarily evidence that something is synthetic.**

The possibility of fabrication changes the meaning of evidence

People have always been able to dispute photographs, recordings and documents.

A photograph could be staged. An audio recording could be edited. Video could be selectively presented. Documents could be forged. Witnesses could disagree about what happened. Digital manipulation expanded these possibilities long before contemporary generative AI.

What changes when synthetic production becomes more accessible is the plausibility of an additional explanation.

Someone confronted with an inconvenient image can say that it was generated. A person hearing an unexpected recording can suggest that the voice was synthesised. An unusual video can be attributed to artificial intelligence even when no specific evidence of generation has been established.

The existence of the technology makes the explanation conceivable.

But **conceivability is not the same as evidence**.

The fact that AI could potentially generate something similar does not establish that AI generated the particular content being encountered. This distinction is easy to overlook because technical possibility can influence intuitive judgement. Once people know that realistic synthetic content exists, the mere realism of a photograph, recording or video may no longer carry the same psychological weight it once did.

This does not eliminate evidential value. It changes one of the assumptions surrounding it.

Doubt can become easier to introduce

Creating uncertainty does not always require persuading someone that an alternative account is true. Sometimes it is enough to make certainty feel less available.

AI-generated content can contribute to this condition because its existence provides another possible explanation for apparently compelling material.

Imagine a genuine recording of an event. Before widespread awareness of synthetic media, disagreement might focus on what the recording shows, what happened before or afterwards, whether it has been edited, or how it should be interpreted. In an environment of realistic generation, another question can be introduced earlier: was the recording itself generated?

That question may be entirely reasonable in some circumstances.

The difficulty arises when the question itself is treated as though it answered what it asks.

Suspicion of synthetic origin and evidence of synthetic origin are different things.

AIGuide 10 does not attempt to provide the technical procedures for resolving such disputes. Nor does it undertake detailed factual verification, which belongs primarily to AIGuide 9. The concern here is the environmental consequence: the growing availability of synthetic media can expand the range of explanations people consider when encountering genuine material.

Genuine content may carry its own burden of uncertainty

Historically, people producing genuine content did not generally need to establish that it had not been generated by AI. That possibility did not exist in its contemporary form.

As synthetic capabilities become more familiar, genuine content can enter an environment in which its origin is no longer automatically assumed from appearance alone.

This is particularly visible with photographs, audio and video because those forms have often been associated with physical recording. But the same issue can affect text.

A student, employee, author or correspondent may write something personally and nevertheless be told that it “sounds like AI”. Such judgements can arise from characteristics such as structure, vocabulary or tone, even though those characteristics are not exclusive to machine-generated writing.

The wider point is not whether any particular suspicion is justified.

It is that **human-created content can increasingly be interpreted through expectations formed by machine-generated content.**

Synthetic media therefore changes the comparative environment in which genuine material is encountered.

Provenance becomes relevant in a new way

Chapter 5 examined provenance as information about where content came from and what happened to it during production and circulation.

The possibility of genuine content being doubted gives provenance another significance.

Information about origin may not only help audiences understand synthetic or modified content. It can also provide context for material that is genuinely recorded or human-created.

This does not mean that provenance automatically proves authenticity. Information about provenance can itself be incomplete, detached or disputed. Different technical and institutional systems may provide different kinds of information, and ordinary users may not always have access to a complete production history.

Nevertheless, the conceptual shift is important.

In an environment containing both convincing synthetic material and genuine content that may resemble it, information about origin can operate in both directions. It can help explain that AI was involved, but it can also help explain that particular content emerged through a different production process.

The issue is therefore not simply detecting artificiality.

It is understanding provenance across a mixed content environment.

Authenticity can become something that is asserted

As uncertainty increases, creators and publishers may sometimes place greater emphasis on describing material as authentic, original, unmodified or human-created.

Such descriptions are themselves part of the changing information environment.

Historically, many forms of content did not require explicit statements about how they were produced. A photograph could simply appear as a photograph. A voice recording could simply be presented as a recording. A written article could carry an author's name without explaining whether automated tools contributed to its composition.

As production methods diversify, those previously implicit assumptions may sometimes become explicit claims.

This does not mean that every piece of content will require a declaration of origin. Nor does it establish what forms of disclosure should be used. Practices may differ among industries, platforms, jurisdictions and communities and may change as technologies and expectations develop.

The broader observation is that **production history can become part of the information communicated about content**, rather than remaining invisible in the background.

Doubt can be appropriate without becoming universal

Recognising that genuine material can be questioned does not imply that questioning is inherently undesirable.

Critical interpretation has always been important. Images can mislead. Recordings can be edited. Sources can be mistaken. Context can be missing. The ability to question what appears convincing can protect against premature conclusions.

The difficulty lies in the opposite extreme: treating the existence of synthetic media as a reason to assume that any inconvenient, surprising or unfamiliar content might therefore be fabricated.

Universal suspicion would not solve the problem created by synthetic content. It would simply replace automatic acceptance with automatic doubt.

Both responses reduce complex questions to a default assumption.

The changing environment instead reveals why certainty can depend upon more than immediate appearance. Sometimes context will make origin clear. Sometimes provenance information will be available. Sometimes uncertainty will remain.

The human challenge is not to eliminate uncertainty from every encounter but to recognise that uncertainty itself can have different foundations.

Social disagreement can widen around origin

When people disagree about content, they may increasingly disagree not only about its meaning but about what kind of object it is.

One person may treat a video as genuine footage while another believes it was generated. One may hear a recording as evidence that someone spoke while another regards the voice as synthetic. One may see a photograph as documentation while another suspects substantial modification.

This creates a different kind of disagreement from disputes about interpretation.

People may no longer share the same starting object.

If one participant believes an image records an event and another believes the image was generated, discussion about what the image means can become secondary to disagreement about its origin.

At a broader social level, such disagreements can complicate shared informational reference points. This does not mean that common understanding becomes impossible. Institutions, communities and individuals continue to use multiple forms of contextual information to establish what they are examining.

But synthetic capability adds another possible layer of disagreement.

The question can shift from “What does this show?” to “Is there anything here that was actually recorded?”

The problem is not distributed equally

The consequences of doubt are also unlikely to be identical across people, societies or contexts.

Access to technical information, trusted institutions, reliable communication channels and contextual knowledge varies. Familiarity with generative technology differs across generations and communities. Some audiences may be highly aware of synthetic media; others may encounter it without understanding how it is produced.

The importance attached to photographs, recordings and institutional sources can also vary culturally and socially.

Different environments may therefore experience the relationship between synthetic capability and genuine content differently.

This reinforces the need for a global rather than technology-centred perspective. There is no single universal level of confidence or suspicion that characterises everyone living with AI-generated content.

What is broadly shared is the underlying possibility: content can now be questioned on the basis of production methods that were previously unavailable or far less accessible.

A changed relationship between seeing and believing

The familiar phrase “seeing is believing” has never been literally reliable. Human perception can be mistaken, images can be misleading and context can transform meaning.

Yet the phrase reflects something important about how visual and recorded evidence has often functioned in everyday life. Direct-looking representations can carry intuitive persuasive force.

Synthetic media complicates that relationship in two directions.

Something can look or sound convincing without having been recorded as it appears.

And something that was genuinely recorded can be doubted because convincing synthetic alternatives are known to exist.

The second consequence is just as important as the first.

It means that the informational effect of AI-generated content cannot be measured simply by counting synthetic items or identifying instances in which people mistake generated material for genuine content. Its existence can also alter the interpretive environment surrounding material that remains entirely genuine.

This creates a deeper form of uncertainty.

The question is no longer only whether synthetic content can imitate reality. It is also whether the possibility of imitation changes how reality itself is received when it is represented through media.

That possibility leads naturally to the next stage of the Guide. When people repeatedly encounter claims, images, voices and narratives within an environment containing both human-created and generated material, familiarity itself can begin to influence perception.

What appears often may feel common. What is repeated may feel familiar. What feels familiar may sometimes acquire an appearance of credibility independently of how many genuinely independent sources lie behind it.

Understanding that relationship between **repetition, familiarity and perceived credibility** is the territory to which the Guide turns next.

Chapter 10 — Familiarity, Repetition and the Appearance of Credibility

People rarely encounter information in isolation. A claim may appear in several places. An image may circulate repeatedly. Similar explanations may emerge across different channels. A phrase may become familiar because it is seen or heard many times.

Repetition is not new to the information environment. Advertising has long relied on repeated exposure. Political messages, public-information campaigns, cultural narratives and commercial slogans can become familiar through recurrence. News organisations may report the same event. Social networks allow material to be copied and redistributed rapidly. Search systems can present multiple pages discussing similar subjects.

Generative AI adds another dimension because repetition no longer needs to consist of identical copies.

The same underlying idea can be expressed in many different words. Similar images can be generated rather than duplicated. A message can appear through different voices, formats and styles. One piece of information can become text, audio, imagery and video. Content can be adapted for different audiences while retaining substantially the same underlying proposition.

As Chapter 7 examined, this capacity can increase the volume and variation of material in circulation. The consequence for human interpretation is subtle but important: **many apparently separate pieces of content do not necessarily represent many independent origins.**

This matters because people can reasonably use recurrence as one of many informal signals when trying to understand the world around them. If an idea appears repeatedly, it may seem widespread. If similar accounts come from apparently different places, they may appear mutually reinforcing. When generation makes variation easier, however, apparent multiplicity and genuine independence can become increasingly difficult to distinguish from appearance alone.

Familiarity is part of ordinary human interpretation

Human beings necessarily rely on patterns.

No person can investigate every statement from first principles. Much everyday understanding develops through accumulated experience. Familiar people, organisations, places, ideas and forms of communication become easier to recognise because they have been encountered before.

Repetition contributes to this familiarity.

A term that initially seems unusual can become ordinary after repeated exposure. A new technology can become conceptually familiar before someone ever uses it. A public issue can

appear increasingly prominent because it is discussed frequently. Repeated imagery can establish visual associations between subjects.

None of this means that familiarity automatically causes belief. People can repeatedly encounter something and continue to reject it. They may even become more sceptical through repetition.

The narrower point is that **familiarity can influence how information is perceived.**

Something encountered for the tenth time does not necessarily feel the same as something encountered for the first time. Repeated exposure can make material easier to recognise, recall and place within an existing mental framework.

Generative AI enters this ordinary human process by changing the potential scale and variety of repeated material.

Repetition can look like independence

Consider the difference between ten genuinely independent accounts and ten reformulations derived from one underlying source.

At the surface level, both may appear as ten pieces of content.

The first may reflect separate observations, research processes, authors or institutions. The second may ultimately trace back to the same material, even though the wording, presentation and format differ.

This distinction predates generative AI. News reports have long drawn from the same agency dispatches. Websites have republished material from common sources. Social-media users can repeat claims originating from a single post. Search results may lead to multiple pages that ultimately rely on the same underlying information.

Generative technology can make this phenomenon less visually obvious because repetition can take the form of variation rather than duplication.

One source can be summarised several ways. Its language can be changed. A claim can be incorporated into different articles. Text can be converted into audio. Similar imagery can accompany different presentations.

The resulting items may look independent even when their informational lineage is closely connected.

This does not make the content inaccurate. A hundred versions of an accurate statement remain capable of being accurate. The issue is different: **the number of versions is not necessarily evidence about the number of independent foundations behind them.**

That distinction belongs to understanding the content environment rather than determining the truth of a particular claim.

Apparent prevalence can be manufactured or emerge naturally

Large quantities of similar content can arise for very different reasons.

In some circumstances, repetition may be deliberate. An organisation, campaign or individual may intentionally distribute many versions of a message in order to increase its visibility.

In other circumstances, no central coordination is required.

Many people may independently ask AI systems about the same popular subject and produce similar material. Organisations may use comparable tools to prepare routine communications. Content creators may respond to the same event. Automated processes may generate variations for different audiences or channels.

The resulting environment can contain substantial repetition even though nobody intended to create a misleading impression of consensus or prevalence.

This distinction is important because environmental effects can arise from ordinary behaviour.

If generation becomes widely used, the information environment may contain more reformulation simply because reformulation has become easier.

The effect experienced by the audience—the repeated encounter with similar ideas—does not necessarily reveal why the repetition occurred.

Visibility is not the same as importance

Repeated exposure can also influence perceptions of what deserves attention.

An issue that appears across many channels can seem socially significant. A product repeatedly encountered in advertising may seem prominent. A phrase that appears throughout online discussion can begin to feel representative of broader public conversation.

Sometimes visibility reflects genuine importance. Major events naturally generate extensive coverage. Widely shared concerns attract repeated discussion. Popular cultural material circulates because many people are genuinely interested in it.

But visibility and importance are not identical.

Digital environments already demonstrate this distinction. Recommendation systems, advertising, search ranking, network effects and viral distribution can make some material far more visible than other material.

Generative AI does not control these distribution systems simply by producing content. As Chapter 7 emphasised, production and distribution are separate.

Yet scalable generation can increase the supply of material available for those systems to distribute.

This can produce an environment in which the frequency with which something is encountered reflects a combination of genuine interest, production volume, distribution mechanisms, repetition and adaptation.

The person receiving the content ordinarily sees only the result.

Repetition can cross media boundaries

One distinctive feature of contemporary generation is the ease with which an underlying idea can move between forms.

A written statement can become an image containing text. It can become synthetic narration. It can form the basis of a short video. It can be summarised, expanded, translated or reformulated for another audience.

The same underlying proposition can therefore appear to a person through several sensory and informational routes.

This cross-media repetition may feel different from seeing an identical post several times. Each encounter can appear fresh because the presentation changes.

A reader may first encounter an idea in text, later hear something similar in audio and eventually see it incorporated into video. Unless the production histories are visible, there may be no obvious indication that the items share a common informational origin.

Again, the significance is not automatically negative.

Cross-media communication can improve accessibility. Translation can broaden participation. Visualisation can make difficult concepts easier to understand. Audio can make information available in circumstances where reading is impractical.

The same capability that increases accessibility can also increase the apparent diversity of repeated material.

Both consequences can exist simultaneously.

Credibility is more than familiarity

Repeated exposure may contribute to an appearance of credibility, but credibility itself is a much broader judgement.

People consider sources, experience, context, consistency, evidence, expertise and many other factors when deciding what weight to give information. Different individuals and cultures may place different emphasis on those factors.

AIGuide 10 therefore does not suggest that repetition determines trust or belief.

Trust in AI and its outputs belongs principally to AIGuide 7. Detailed factual verification belongs principally to AIGuide 9.

The narrower concern here is how the structure of a synthetic information environment can affect **the cues surrounding content before those deeper judgements are made**.

If many apparently distinct items express similar ideas, that multiplicity can form part of the context in which the ideas are encountered. Understanding whether those items represent independent origins or repeated derivations may therefore become relevant to understanding what their recurrence signifies.

This is another application of provenance.

Provenance is not only about whether a photograph was generated or whether a voice was synthetic. It can also concern informational lineage: whether apparently separate material ultimately derives from common sources or production processes.

Human attention remains the limiting resource

The scale of content also interacts with a constraint that generation cannot remove: human attention.

People have limited time to investigate what they encounter. When information is abundant, they necessarily select what to read, watch, hear and remember.

Repeated content can occupy a larger share of that limited attention simply because it appears more often.

This does not guarantee persuasion. Repetition can annoy, bore or alienate audiences. Excessive volume can reduce rather than increase attention. People develop ways of ignoring material they regard as repetitive or irrelevant.

Nevertheless, an environment in which some ideas can be represented in large numbers of variations raises a structural question about how finite human attention is distributed across an expanding supply of content.

The answer depends on many factors beyond AI: platform design, social networks, institutional choices, personal interests, language, geography and culture among them.

Generative AI is one contributor to this broader environment, not its sole cause.

Global information environments will differ

The effects of repetition and familiarity will not be uniform around the world.

People operate within different media systems, languages, levels of connectivity and cultural contexts. Some information environments contain highly diverse sources; others may be more concentrated. Access to generative technologies differs, as do the platforms through which generated content circulates.

Language is particularly relevant.

Generation and translation can increase the amount of material available in languages or contexts where producing localised content previously required greater resources. This may broaden access and representation.

At the same time, scalable translation and adaptation can allow the same underlying narratives to travel across linguistic boundaries in many forms.

Neither development has one predetermined consequence.

They illustrate why the effects of generation at scale must be understood within the informational environments in which people actually live.

Familiarity can become detached from origin

Across the earlier chapters, a recurring theme has emerged: surface characteristics increasingly reveal less about production history.

Repetition extends this principle from individual content to collections of content.

Ten different-looking items may originate independently.

They may also share a common source.

They may be human-created, AI-assisted, substantially generated or mixed.

Their similarity may arise deliberately, accidentally or because many creators are responding to the same underlying event.

The audience may not know which explanation applies.

This does not mean that repeated information should automatically be discounted. Such a conclusion would be as simplistic as assuming that repetition proves credibility.

Instead, the changing environment makes another distinction increasingly important:

frequency of appearance is not necessarily frequency of independent origin.

That distinction matters because people understand the world partly through what appears around them. When scalable generation can multiply representations, reformulations and formats, apparent prevalence can become increasingly separated from the number of underlying informational sources.

The consequence extends beyond individual judgement.

Shared informational environments are built from the material that communities encounter collectively. If people experience different combinations of repeated, generated, modified and

human-created content, their sense of what is common, significant or widely accepted may also differ.

The next chapter therefore moves from individual familiarity to a broader social question: **what happens to shared reality when communities increasingly inhabit an environment containing synthetic content whose origins, purposes and relationships with other material may not always be visible?**

Chapter 11 — Shared Reality in an Environment of Synthetic Content

People do not understand the world only as individuals. Much of what is known about events, institutions, communities and other people comes through information that is encountered collectively. News reports, photographs, recordings, public statements, educational material, personal communication and shared cultural references help groups develop some common understanding of what has occurred and what is happening around them.

That understanding has never been perfectly uniform.

People interpret the same event differently. Communities draw on different sources. Cultural, linguistic and historical experiences shape perception. Political and social disagreement existed long before digital media or artificial intelligence. Even when people agree that an event occurred, they may disagree profoundly about its causes, significance or consequences.

A shared reality therefore does not require everyone to hold the same opinions.

At a more basic level, it depends upon the possibility of establishing enough common reference points for people to understand what they are discussing. A photograph is recognised as a photograph of a particular event. A statement is understood to have come from a particular person or institution. A recording is accepted as representing something that was said. People may then disagree about what those materials mean.

Synthetic content introduces another potential layer of disagreement.

If one person believes an image records an event while another believes it was generated, they are not merely interpreting the same evidence differently. They may disagree about **what kind of evidence exists in the first place**.

Shared reality has always been mediated

Most people directly witness only a small fraction of the events that contribute to their understanding of the wider world.

They rely on forms of mediation.

Journalists report events that readers did not attend. Photographs show places viewers have never visited. Recordings preserve words spoken elsewhere. Educational materials explain historical events that occurred before anyone now living could observe them. Scientific instruments represent phenomena beyond ordinary human perception.

Modern societies therefore depend heavily on mediated knowledge.

This is not inherently a weakness. Shared information allows people to understand circumstances far beyond their immediate experience.

But mediation creates relationships among events, representations, sources and audiences. People ordinarily make assumptions about those relationships. A news photograph is expected to have some connection with the event it accompanies. A quotation attributed to a public official is understood to represent words the official actually communicated. Documentary footage is approached differently from fictional reconstruction.

As earlier chapters have shown, generative AI can complicate these relationships by making realistic representation increasingly separable from physical occurrence.

The social consequence is not that mediated knowledge ceases to function. It is that **the production status of the material supporting shared understanding may itself become another subject of disagreement.**

Communities need not see the same information

Digital information environments were already fragmented before generative AI became widespread.

Different people follow different news sources, social networks, creators and communities. Search results can differ according to language, location and context. Recommendation systems can expose users to different material. Personal relationships influence what is shared through private messaging.

As a result, two people living in the same city can experience substantially different streams of information about the same wider world.

Generated content enters this already diverse environment.

Its capacity for variation can make material easier to adapt across languages, formats and audiences. This can broaden communication and make information accessible to people who might otherwise receive less relevant material.

At the same time, scalable variation means that different groups may encounter different representations of the same underlying subject.

One audience may receive a concise text explanation. Another may encounter synthetic narration. A third may see generated imagery or simulation. Content may be adapted to cultural or linguistic contexts.

Variation does not necessarily produce disagreement. Different representations can communicate substantially the same underlying information.

But the more varied the informational pathways become, the less it can be assumed that people discussing a subject have encountered the same material or understood its production in the same way.

Disagreement about origin differs from disagreement about meaning

Consider two people watching the same video.

Both may agree that the video exists. They may disagree about whether the actions shown were appropriate, what motivated the people involved or what consequences should follow. Those are disagreements about interpretation.

Now consider a different situation.

One person believes the video records an event. The other believes the scene was generated.

The disagreement occurs at an earlier stage.

Before they can discuss what the event means, they may first need a common understanding of whether the represented event occurred in the form shown.

This distinction illustrates why synthetic content can affect shared reality without needing to persuade everyone of a particular falsehood.

Uncertainty about origin can itself complicate the establishment of common reference points.

Chapter 9 examined the corresponding problem for genuine content: once convincing synthetic media exists, authentic material can also be questioned. At a social level, the combination of these possibilities can create situations in which different groups assign different production histories to the same content.

The question becomes not merely “What should we think about this?” but “What exactly are we looking at?”

Shared reality does not require automatic trust

It would be misleading to suggest that societies previously operated through universal agreement about the authenticity of information.

Propaganda, forgery, misleading imagery, fabricated documents and contested accounts have long existed. Institutions have made mistakes. Sources have disagreed. People have questioned official narratives and challenged supposed evidence.

Critical scrutiny is therefore not an obstacle to shared reality. It can be part of how shared understanding develops.

The distinction lies between **questioning a claim** and **losing the possibility of establishing common informational reference points altogether**.

Generative AI does not make the latter inevitable.

Many forms of context remain available. Sources can provide information about production. Events may have multiple independent records. People can directly witness circumstances. Institutions and communities can develop conventions surrounding generated material. Provenance can sometimes clarify how content was created.

The existence of synthetic content therefore does not imply that societies must descend into universal uncertainty.

It does mean that shared understanding may increasingly depend on information surrounding content rather than solely on the apparent realism of the content itself.

Different institutions play different informational roles

Shared reality is supported through many institutions and social practices.

Journalism, education, science, government, courts, archives, libraries, businesses, professional bodies and community organisations all create and preserve information in different ways. Families and personal networks also establish their own forms of shared knowledge.

Generative AI can enter each of these environments differently.

An educational institution may use generated illustrations. A business may use synthetic narration. A media organisation may incorporate generated visualisations. A public body may use machine-assisted translation. Creative industries may integrate synthetic elements extensively.

The mere presence of AI does not determine whether these institutions remain credible or whether their information is accurate.

What changes is that organisations may increasingly produce content through processes that audiences cannot infer simply by examining the finished material.

Institutional context can therefore become part of how people understand content's relationship with reality.

This observation should not be converted into a rule that particular institutions should automatically be trusted. Trust is contextual and is examined primarily elsewhere in the AIGuide series.

The narrower point is that shared informational environments depend on more than individual pieces of media. They also depend on relationships among sources, audiences, practices and expectations.

Synthetic content can create shared experiences too

The relationship between generated content and shared reality is not exclusively disruptive.

Synthetic material can itself become part of shared cultural experience.

People may collectively watch generated entertainment, use synthetic characters, engage with virtual environments or encounter AI-assisted creative works. Educational simulations can provide groups with common representations of concepts that are difficult to observe directly. Generated translations may allow material to reach communities across language boundaries.

In these circumstances, people can share an experience while fully understanding that its content is synthetic.

A fictional film does not need to document a real event in order to become culturally meaningful. Similarly, generated material does not need to be mistaken for recorded reality in order to contribute to common experience.

This distinction is important.

Shared reality is not composed only of literal records of physical events. Human societies have always shared stories, symbols, art and imagined worlds.

The informational challenge arises when **different relationships with reality are confused**, not simply when synthetic content exists.

Scale can affect the appearance of collective agreement

Chapter 10 examined how repetition and variation can create apparent multiplicity without necessarily reflecting multiple independent origins.

At the social level, this can influence perceptions of what other people believe or what ideas are widely held.

If similar messages appear repeatedly across many forms, they may create an impression of broad prevalence. That impression may sometimes accurately reflect widespread interest. In other cases, the visible volume of content may exceed the number of independent people or sources behind it.

Generated variation makes the distinction potentially harder to infer from appearance.

This does not mean that apparent online consensus is necessarily artificial. Nor is generative AI required for perceptions of consensus to become distorted. Coordinated communication, advertising, automated accounts and repeated sharing all predate contemporary generative systems.

AI adds capacity for producing diverse-looking material at scale.

The durable distinction is therefore between **the amount of content expressing a position and the number or independence of the people, sources or processes represented by that content**.

The two may correspond, but they need not.

Shared reality is also culturally situated

There is no single global information environment.

Languages, media institutions, communication practices, levels of digital access and social expectations differ widely. Communities may have different traditions regarding authority, evidence, privacy, public communication and the representation of individuals.

Synthetic content will enter these environments unevenly.

Some populations may encounter extensive generated material through commercial platforms. Others may experience it primarily through translation, entertainment or institutional communication. Technical access and media literacy vary. So do regulatory responses and cultural expectations about disclosure.

A globally applicable analysis therefore cannot assume one universal pathway through which synthetic content affects social understanding.

What can be identified is the common underlying development: more communities will increasingly need to interpret information in environments where realistic content may have a wider range of possible production histories.

How those communities respond will depend on context.

Common ground in a more complex environment

Shared reality has never meant complete agreement.

People can disagree about values, explanations and conclusions while still sharing enough common ground to conduct meaningful discussion. Indeed, disagreement often depends upon having something identifiable to disagree about.

AI-generated content introduces new complexity into the formation of that common ground because representation, origin and occurrence can increasingly be separated.

An image can exist without the depicted scene having been photographed. A familiar voice can exist without the recognised person speaking the words. A realistic event can be simulated. Genuine material can be doubted because synthetic alternatives are technically possible. Repeated variations can create apparent prevalence without revealing how many independent origins lie behind them.

None of these developments makes common understanding impossible.

They do, however, make the informational environment more layered.

People may increasingly need to distinguish among what content depicts, how it was produced, where it came from, what purpose it serves and what relationship it has with the reality it appears to represent.

Those distinctions will not carry equal significance in every situation.

A synthetic scene in entertainment, a generated illustration in education, an altered image in advertising, a simulated event in training and a photograph presented as documentation each operate within different contexts.

Understanding those differences is therefore essential to avoiding an overly general account of AI-generated content.

The next chapter turns directly to that variation.

Rather than asking what synthetic content means in the abstract, it examines why **the same underlying capabilities can carry very different consequences in journalism, entertainment, education, advertising, creative work, workplace communication and personal communication.**

The technology may be similar.

The context can change almost everything about how the content is understood.

Chapter 12 — Different Contexts, Different Consequences

AI-generated content does not have the same meaning everywhere it appears.

A synthetic image used in a fictional film occupies a different informational position from a similar image presented in a news report. An AI-generated explanation used in an educational setting raises different questions from generated text appearing in an advertisement. A synthetic voice used openly for narration differs from the same kind of voice appearing in a personal message that seems to come directly from someone known to the listener.

The underlying technologies may be similar. The content may even look or sound similar.

The context changes what audiences reasonably understand the content to represent.

This is why the consequences of AI-generated content cannot be assessed solely by asking whether AI was involved. The earlier chapters have established distinctions among origin, production, authenticity, provenance, representation and factual truth. Context brings those distinctions together by determining which of them matters most in a particular setting.

There is no single universal expectation that applies equally to journalism, entertainment, education, advertising, creative work, workplace communication and personal communication. Each operates according to different purposes and conventions. Those conventions also vary across countries, cultures, organisations and communities.

Understanding the emerging content environment therefore requires attention not only to **what content is**, but also to **where it appears and what role it is expected to perform**.

Journalism and the expectation of representation

Journalism provides a particularly clear example of why context matters.

News and current-affairs content can include photographs, video, graphics, illustrations, reconstructions, maps and other forms of representation. These forms have never all carried the same meaning. An illustration accompanying an article is not necessarily understood as photographic evidence, while footage presented as recording an event ordinarily carries a closer relationship with what physically occurred.

Generative AI expands the range of material that can appear within this environment.

Synthetic imagery may illustrate events for which no photograph exists. Generated visualisations may help explain complex subjects. Machine-assisted translation can broaden access to reporting. AI may also contribute to text production, transcription or other parts of the publishing process.

The significance of these uses depends on what audiences understand them to represent.

A generated illustration that is understood as an illustration occupies a different position from a generated image that appears to document an event. The distinction does not arise because one is inherently acceptable and the other inherently unacceptable. It arises because journalism carries particular expectations about the relationship between representation and occurrence.

Detailed questions about whether individual journalistic claims are factually correct belong primarily to AIGuide 9. The concern here is the more fundamental distinction between **content that represents information and content that appears to record the reality being reported**.

Entertainment operates through constructed realities

Entertainment provides a contrasting context.

Film, television, games, animation and other creative forms have always constructed experiences. Actors portray people they are not. Sets represent other places. Editing changes time and sequence. Visual effects create events that never occurred.

Audiences ordinarily enter such environments with some understanding that what they see may be constructed.

Generative AI can become another component of that construction. Backgrounds may be generated. Fictional characters may be synthetic. Voices, imagery or sequences may be created or modified computationally.

In this context, the fact that something did not physically occur before a camera may be entirely consistent with the purpose of the content.

Authenticity may therefore concern something other than documentary origin. Audiences may be interested in artistic intention, performance, attribution or the relationship between human and machine contribution.

Those questions can be important, but they demonstrate why “Is this real?” is sometimes the wrong question.

A fictional dragon is not misleading merely because no dragon stood before a camera.

The relevant expectation is different.

Education combines explanation and representation

Educational content occupies another distinctive position because its purpose often involves helping people understand things they cannot directly observe.

Diagrams simplify reality. Historical reconstructions depict periods for which no visual recording exists. Models represent processes. Hypothetical scenarios allow learners to explore consequences. Illustrations make abstract ideas more accessible.

Generated content can extend these established practices.

An educator may use synthetic imagery to visualise an historical environment, generated narration to make material available in another form, or machine-assisted text to explain a concept at different levels of complexity.

The generated nature of the material does not determine its educational value.

At the same time, educational contexts can make distinctions between **illustration and evidence** particularly important. A generated reconstruction of an ancient settlement may help learners imagine its possible appearance, but it is not thereby a photographic record of that settlement. A synthetic depiction of a scientific process may clarify an idea while remaining a representation rather than direct observation.

The consequence depends on what the content is understood to show.

AI-generated educational material therefore illustrates a broader principle: representation can be useful without being mistaken for recording.

Advertising routinely combines information and imagination

Advertising has long occupied the boundary between factual information, persuasion and imaginative representation.

Products may be photographed in idealised settings. Visual effects can create scenes designed to attract attention. Actors may represent customers. Images may communicate mood or aspiration rather than literal descriptions of ordinary experience.

Generative AI expands the range and scale of such content.

Images can be adapted for different audiences. Synthetic people or settings can appear in campaigns. Text can be reformulated. Product concepts can be visualised before physical versions exist.

The informational significance of these practices depends on what audiences understand the content to communicate.

A clearly imaginative background may function primarily as visual expression. A synthetic representation that appears to show a product characteristic may carry a different relationship with consumer understanding.

The Guide does not attempt to establish legal or regulatory boundaries for advertising, which differ across jurisdictions. Nor does it prescribe how organisations should use generated material.

The relevant observation is that advertising already blends representation and persuasion, making context essential to interpreting what synthetic content appears to claim.

Creative work brings process into the foreground

Creative contexts raise another set of questions because the process of creation may itself contribute to how a work is understood.

A novel, painting, song, photograph or film can be valued not only for the finished object but also for its relationship with human imagination, experience, skill and expression.

Generative AI complicates this relationship without resolving it in one direction.

Some creative works may involve limited AI assistance. Others may combine extensive human direction with generated elements. Some may be substantially generated and then selected or modified by people. Chapter 2 described this wider continuum.

Different audiences and creative communities may attach different significance to these distinctions.

For one viewer, the finished image may be the primary consideration. For another, knowing how it was made may substantially affect its meaning. Cultural expectations about originality, authorship and creative labour can also differ and evolve.

AIGuide 10 does not determine what should count as art or who should be recognised as an author. Its concern is that **production history can become part of the interpretive context of creative content.**

The same image may be experienced differently once its production process is known.

Workplace communication can make authorship ambiguous

AI-generated content also enters routine organisational life.

Reports, presentations, correspondence, summaries, translations, meeting materials and other communications can involve AI at different stages.

In many circumstances, the precise wording process may be less significant than the fact that an individual or organisation stands behind the final communication. A manager may use assistance to improve a message while still intending every substantive point it contains. An organisation may use automated systems to produce routine communications according to established information.

In other circumstances, the degree of human involvement may matter more.

A document may appear to reflect someone's analysis when much of that analysis was generated. A personalised message may feel less personal if it was substantially automated. Different professional settings may attach different expectations to disclosure, authorship and review.

These questions cannot be resolved through a universal rule because workplaces differ enormously.

The broader environmental change is that receiving a polished document from a colleague no longer necessarily reveals how much of its formulation occurred through direct human composition.

As elsewhere, **apparent authorship and production history can diverge.**

Personal communication carries expectations of presence

Perhaps one of the most humanly immediate contexts is personal communication.

Messages between family members, friends and acquaintances often carry assumptions that differ from institutional communication. A voice message can convey not merely information but presence. A photograph can communicate an experience. Personal wording can express attention, affection, humour or concern.

AI can assist these communications in ordinary ways.

Someone may use translation to communicate across languages. A person may use writing assistance because expressing themselves is difficult. Accessibility technology may help someone communicate through synthetic speech. Generated imagery may be shared playfully or creatively.

Machine involvement therefore does not necessarily make communication less personal.

But context again determines significance.

A message understood to be personally composed may be experienced differently if substantially generated. A familiar voice can carry emotional meaning partly because the listener assumes the person actually spoke. An image shared as a record of someone's experience carries a different implication from an imaginative image created for amusement.

The relevant issue is not whether technology belongs in personal communication.

It is that **personal communication often carries expectations about human presence and intention that extend beyond the literal information being transmitted.**

Synthetic content can affect those expectations.

The same content can change meaning when moved

Context is not permanently attached to content.

A generated historical reconstruction may be clearly identified as such within an educational programme. If extracted and circulated separately, it may appear to be archival footage. An advertising image may be reposted without the campaign that made its imaginative purpose obvious. A fictional clip may circulate as though it recorded an actual event.

The production process has not changed.

The surrounding information has.

This demonstrates why provenance and context are closely related. Understanding where content came from can help restore information that disappears when material moves between environments.

Digital circulation makes such movement routine. Screenshots, excerpts, reposts and short clips can travel far from their original setting.

AI-generated material participates in this broader pattern while adding content whose synthetic nature may not be visible from appearance alone.

No context exists in isolation

The categories considered here also overlap.

Journalism can contain creative illustration. Education can use entertainment. Advertising can resemble personal communication. Workplace materials can circulate publicly. Creative content can become part of news coverage.

The objective is therefore not to establish rigid boundaries among sectors.

It is to recognise that **expectations change with purpose**.

The same synthetic technique can have different informational consequences depending on whether audiences expect documentation, illustration, imagination, simulation, persuasion, personal expression or organisational communication.

Those expectations are neither globally uniform nor permanently fixed.

Different societies may develop different conventions around generated content. Industries may adopt distinct practices. Individuals may become increasingly accustomed to some forms of synthetic media while continuing to attach special significance to others.

Context prevents the technology from becoming the whole explanation

A recurring risk in discussing AI-generated content is to treat the technology itself as though it determines the meaning of everything produced through it.

The preceding examples show why that is insufficient.

A generated image is not one thing informationally.

It can illustrate a lesson, advertise a product, create an imaginary world, simulate a proposed building, represent an organisation or appear to document an event.

Synthetic audio can provide accessibility, perform fiction, communicate institutional information or represent the apparent voice of an individual.

Generated text can explain, advertise, entertain, assist or communicate personally.

The method of production matters, but it is only one part of interpretation.

Purpose, presentation, provenance and context shape what the content means to the person encountering it.

This insight brings the Guide toward its final synthesis.

People cannot always know the full production history of what they encounter. Context can be incomplete. Content can move away from its original setting. Human and machine contributions can be intertwined. Authenticity can have several meanings. Genuine material can be doubted, while generated material can appear entirely ordinary.

The result is not a world in which interpretation becomes impossible.

It is an environment in which interpretation may increasingly take place under conditions of unresolved origin and incomplete context.

The next chapter therefore turns directly to that human condition: **what it means to interpret content when some uncertainty about its creation, authorship or modification cannot readily be removed.**

Chapter 13 — Interpreting Content Under Uncertainty

The preceding chapters have described an information environment in which the origins of content can be increasingly difficult to infer from appearance alone. Human and machine contributions may be intertwined. Text, images, voices and video carry different signals about their production. Authenticity can mean different things in different circumstances. Provenance may be incomplete. Synthetic representations can resemble recorded reality. Content can be produced and varied at scale, and genuine material can itself become easier to question.

These developments lead to a distinctly human condition: **people may sometimes need to interpret content without being able to resolve every uncertainty surrounding how it was created.**

This is not entirely new. Human beings have always made sense of incomplete information. The author of a document may be uncertain. The context of a photograph may be missing. People can disagree about what a recording means. Historical evidence is often fragmentary. Everyday conversations contain ambiguity. Decisions about what something signifies routinely occur before every relevant fact is known.

AI-generated content adds new forms of uncertainty to this familiar condition.

A person may not know whether an image began with a camera. A voice may sound recognisable without revealing whether the recognised person actually spoke the words. A passage may appear under someone's name without showing the extent of machine assistance behind it. Video may combine recorded and generated elements in ways that are not visible from the finished sequence.

In some circumstances, these uncertainties matter greatly. In others, they may have little bearing on what the content is being used for.

The challenge is therefore not simply uncertainty itself. It is understanding **what kind of uncertainty exists and whether it is significant to the interpretation being made.**

Not every uncertainty has the same importance

Consider an illustration accompanying a fictional story.

A reader may not know whether it was drawn manually, produced digitally, generated through AI or created through some combination of methods. If the image is understood simply as an imaginative representation, uncertainty about its production may have little effect on the story being communicated.

Now consider an apparently photographic image presented as recording a particular public event.

Whether that image began with an actual camera may be much more significant to what viewers understand it to represent.

The uncertainty is similar in one respect: the production process is not known.

Its consequence is different because the context is different.

This distinction prevents uncertainty from becoming a universal reason for suspicion. The fact that something is unknown does not establish that the unknown feature matters equally in every setting.

Meaning depends partly on what the content claims or appears to be doing.

Several uncertainties can coexist

AI-generated content can create uncertainty across several dimensions at once.

There may be uncertainty about **origin**: whether material was recorded, generated or both.

There may be uncertainty about **authorship**: who developed the ideas, composed the words or selected the final material.

There may be uncertainty about **modification**: whether an original item has been changed and, if so, how significantly.

There may be uncertainty about **representation**: whether depicted people, places or events correspond to physical occurrences.

There may also be uncertainty about **context**: where the content first appeared, why it was created or what audiences were originally expected to understand.

These questions are related, but they are not interchangeable.

Knowing that an image was AI-modified does not necessarily explain whether the modification changed something important. Knowing that text was generated does not determine whether the information it communicates is factually correct. Knowing that a voice is synthetic does not establish whether the person whose voice it resembles authorised the message.

This is why interpretation becomes more precise when different uncertainties are kept conceptually separate.

Otherwise, one unknown can easily be treated as though it answers another question.

Uncertainty is not evidence of artificiality

One of the most important consequences of the mixed content environment is that uncertainty operates in both directions.

A person may be unable to establish that something is human-created.

But inability to establish human creation does not establish AI generation.

Similarly, uncertainty about whether an image was recorded does not mean it was generated. Suspicion that a voice may be synthetic does not establish that it is. A writing style that resembles machine-generated language does not demonstrate its production history.

Chapter 9 examined how genuine content can become easier to doubt once people know that convincing synthetic alternatives exist.

That development creates a temptation to interpret unresolved origin as evidence of synthetic origin.

The two are not the same.

Sometimes the most accurate understanding available to an ordinary observer may simply be that the production history is unknown.

That position may feel less satisfying than classification, but it preserves an important distinction between **what is known and what is inferred**.

Interpretation begins before verification

People encounter far more content than they can formally investigate.

A person scrolling through information, reading routine correspondence, viewing advertising or watching entertainment does not normally conduct a detailed verification process for every item encountered. Much interpretation occurs rapidly and informally.

This was true before generative AI and remains true now.

People draw on source, context, experience, familiarity and expectations. They recognise conventions. They distinguish, often effortlessly, between a fictional film and a news broadcast, between an advertisement and a personal message, between an illustration and a family photograph.

Generated content enters these existing interpretive processes.

The question for AIGuide 10 is therefore not how every piece of content should be verified. Detailed factual verification belongs primarily to AIGuide 9.

The question is what changes in ordinary interpretation when some of the signals people historically relied upon become less conclusive.

One answer emerging across this Guide is that **surface appearance may increasingly need to coexist with contextual understanding rather than substitute for it**.

This is an observation about the environment, not a prescribed procedure for navigating it.

Context can reduce uncertainty without eliminating it

Chapter 12 demonstrated how the significance of synthetic content changes across journalism, entertainment, education, advertising, creative work, workplaces and personal communication.

Those contexts also shape uncertainty.

A generated image inside an obviously fictional work may be easy to interpret even if its precise production history is unknown. A synthetic reconstruction may be understood because its educational context explains what it represents. An institutional communication may be meaningful because the organisation clearly stands behind it even if AI contributed to its wording.

Context can therefore make content understandable without providing a complete technical history.

But context is not always preserved.

Images are reposted. Video clips are extracted. Screenshots remove surrounding information. Text is copied. Audio travels between platforms. As Chapter 5 observed, content can become separated from information about its provenance.

In those situations, uncertainty may increase because the material has lost some of the signals that previously explained it.

The content itself may remain unchanged while the basis for interpreting it has weakened.

Human judgement remains contextual

The AISF AIGuide architecture places humans at the centre of interpretation and decision-making. In the context of generated content, that principle does not mean that individuals are expected to possess perfect technical knowledge or to resolve every ambiguity themselves.

Human judgement is broader than technical detection.

It includes recognising what kind of situation is being encountered, what the content appears to represent, what information about its origin is available and what remains uncertain.

The significance of those questions changes according to circumstance.

An unexplained production history may be relatively unimportant when viewing decorative imagery. It may become more consequential when content appears to represent someone's actions or words.

Human judgement therefore operates not by making all uncertainty disappear but by understanding uncertainty in relation to context and consequence.

This distinction also preserves the boundary with AIGuide 8. That Guide owns the broader territory of AI and human decision-making. AIGuide 10 is concerned more narrowly with the

interpretive conditions surrounding synthetic content before or alongside whatever decisions may later be made.

Confidence and certainty are not identical

People routinely act on understandings that fall short of absolute certainty.

They recognise a friend's voice without conducting technical authentication. They accept that a family photograph depicts a real gathering because they understand its context. They read institutional communications without reconstructing their full production histories.

AI-generated content does not necessarily make such ordinary confidence impossible.

It does, however, demonstrate why confidence and certainty should not be treated as identical concepts.

A person may have substantial contextual reasons for understanding content in a particular way while acknowledging that its complete production history is unknown.

Conversely, content may look completely convincing while its context provides little basis for understanding where it came from.

The strength of visual or auditory realism and the strength of contextual understanding can therefore diverge.

This is another reason the emerging environment cannot be reduced to whether synthetic content becomes technically indistinguishable from recorded content. Human interpretation has never depended only on sensory realism.

People interpret relationships, sources, situations and purposes as well.

Uncertainty can itself become part of communication

There is another consequence of this environment.

The uncertain status of content can become part of what people communicate about it.

A person may share an image while noting that its origin is unclear. A publication may describe material as reconstructed or generated. An organisation may explain that AI contributed to production. Conversely, people may contest such descriptions.

Information about uncertainty therefore becomes another layer surrounding the original content.

This can be useful because it distinguishes what is established from what remains unresolved.

But descriptions of uncertainty can themselves travel separately from their basis. A claim that an image “may be AI-generated” can circulate even when no additional evidence is provided. Repetition can then make the suspicion itself familiar, as Chapter 10 examined.

Uncertainty is therefore not merely an absence of knowledge. In a networked information environment, statements about uncertainty can become content in their own right.

Different people will experience uncertainty differently

There is no universal experience of synthetic content.

People differ in technical familiarity, media experience, language, age, cultural background and access to contextual information. Some may encounter generated media frequently and develop strong expectations about it. Others may encounter it less visibly.

Different societies also operate through different media institutions and communication practices.

As a result, the same item may produce very different levels of uncertainty for different audiences.

One person may immediately recognise that a video is intended as simulation because they understand the context in which it was created. Another encountering an extracted version may interpret it differently. A synthetic voice may be obvious to someone familiar with the technology while sounding entirely ordinary to someone else.

This variability reinforces the limits of assuming that the content itself determines how it will be understood.

Interpretation occurs through the relationship between **content and audience**, not through content alone.

Living with unresolved origin

The central question of AIGuide 10 asks what changes when seeing or hearing something is no longer sufficient reason to assume how it was created.

By this stage, the answer has become clearer.

The change is not that people cease to understand content.

It is that production history can become an additional dimension of interpretation—sometimes visible, sometimes disclosed through provenance, sometimes inferred from context and sometimes unresolved.

That uncertainty does not carry the same significance everywhere.

It does not establish falsehood.

It does not establish synthetic origin.

It does not require every piece of content to be treated as suspicious.

Nor can realistic appearance alone always remove it.

The emerging environment therefore asks more of the distinction between **what content appears to be and what is actually known about how it came into existence.**

Human judgement remains relevant precisely because those two things will not always align perfectly.

This Guide has moved from the weakening connection between appearance and origin through the continuum of human and machine production, different media forms, authenticity, provenance, synthetic representation, scale, simulation, doubt, repetition, shared reality and contextual difference.

One final question remains.

These technologies will continue to change. Today's obvious synthetic characteristics may not remain obvious. Production practices will evolve. Social conventions may develop. New forms of content may emerge.

The durable challenge is therefore not how to understand one fixed generation of AI-generated media.

It is how to understand **an information environment whose relationship with synthetic content will itself keep changing.**

That is the final territory to which the Guide turns.

Chapter 14 — Living with an Information Environment That Keeps Changing

AI-generated content is not a settled category within a settled information environment.

The technologies that produce synthetic text, images, audio and video continue to develop. The ways people use those technologies are changing alongside them. Institutions, creative communities, businesses, educators, media organisations, platforms and individuals are developing new practices. Social expectations about authorship, authenticity and disclosure are evolving. Forms of content that once appeared obviously synthetic can become more realistic, while audiences themselves become more familiar with the possibility of generation and modification.

Any account of AI-generated content that depends too heavily on the characteristics of today's systems therefore risks becoming outdated.

The more durable question is the one that has guided this Guide from the beginning: **what changes when people can no longer assume how content was created simply from seeing, hearing or reading it?**

Across the preceding chapters, that question has revealed a broader transformation. The issue is not merely that machines can generate convincing material. It is that the relationships among content, origin, authenticity, context and human interpretation are becoming more complex.

Living with AI-generated content therefore means living within an information environment that is itself changing.

The technology will change faster than many conventions

New technologies often develop before societies establish settled expectations around them.

Early practices can be inconsistent. Terminology may differ. Institutions may respond in different ways. What initially appears unusual can become routine, while practices once considered normal may acquire new meanings.

Generative AI is developing within this pattern.

The distinction between generated and recorded imagery may evolve as production methods change. Synthetic voices may become increasingly common in some forms of communication. AI assistance in writing may become routine in particular contexts while remaining significant in others. New methods of describing content origin may emerge, and existing approaches may be revised.

These developments are unlikely to occur uniformly around the world.

Different jurisdictions may establish different legal or regulatory approaches. Industries may develop their own practices. Cultural expectations surrounding creativity, authorship, identity

and representation may vary. Technical capabilities may spread unevenly across languages and regions.

There may therefore never be one moment at which the global community arrives at a single settled understanding of synthetic content.

The information environment can instead continue to evolve through many overlapping technological and social changes.

Some distinctions will remain even as technologies change

Although particular technologies may become outdated, several conceptual distinctions examined throughout this Guide are likely to remain useful.

Appearance and origin are not the same thing.

Content can resemble a photograph without necessarily originating from a camera. A voice can resemble a person without necessarily corresponding to that person's act of speaking. Fluent text does not disclose its production history.

Human-created and AI-generated are not always opposing categories.

Content can move through multiple stages of human and machine contribution. Assistance, modification, generation, selection and editing can coexist within one production process.

Authenticity and factual truth are different questions.

Human-created content can be false. Generated content can communicate accurate information. Something can be authentic in relation to its source while containing machine-generated elements, or synthetic in production while openly representing itself as such.

Provenance and appearance provide different kinds of information.

The finished surface of content tells people what they encounter. Its provenance concerns where that content came from and what happened during its production and circulation.

Representation and occurrence can be separated.

A realistic depiction of a person, place or event does not necessarily establish that the depicted circumstances physically occurred in the form presented.

These distinctions do not depend on one generation of AI technology. They describe relationships that remain relevant even as the tools themselves change.

Synthetic content will increasingly coexist with other content

It is tempting to imagine a future information environment divided neatly between human-created and AI-generated material.

The continuum examined in Chapter 2 suggests a more complex reality.

Many forms of content may contain different levels of machine involvement. Some material will remain substantially human-created. Some will be substantially generated. Much may occupy positions between those points.

The result is likely to be coexistence rather than replacement.

Photographs will continue to exist alongside generated imagery. Human voices will coexist with synthetic narration. People will continue to write while also using machine assistance. Recorded video will circulate alongside simulation, generated sequences and mixed productions.

The informational challenge arises partly because these forms can share the same digital spaces.

A social feed can contain a family photograph, a generated advertisement, documentary footage, a fictional clip, a synthetic illustration, a human-written comment and machine-assisted text within moments of one another.

The environment does not necessarily announce a transition between them.

People move continually among different relationships with reality.

Familiar conventions may adapt

Societies have repeatedly developed conventions for distinguishing among forms of representation.

Audiences understand that novels are not news reports. Illustrations are treated differently from photographs. Dramatisations differ from documentaries. Advertising operates according to conventions different from personal communication.

Generative AI may contribute to the development of additional conventions.

Some synthetic content may become so familiar within particular contexts that its production method attracts little attention. Other uses may continue to make origin highly significant. Practices for communicating generation, modification or provenance may become more common in some environments.

The form these conventions take cannot be assumed in advance.

Nor will their development necessarily remove ambiguity. Established media conventions have never prevented all misunderstanding, deception or disagreement.

Their importance lies instead in providing contextual signals through which audiences understand what kind of content they are encountering.

As Chapter 8 showed, information, expression and simulation can use similar technologies while maintaining different relationships with reality.

The continuing development of conventions may help audiences interpret those differences even when surface appearance does not.

Greater realism does not remove the need for context

Discussion of synthetic media often focuses on whether generated content will become indistinguishable from recorded or human-created material.

That question is technologically important, but from a human perspective it captures only part of the issue.

People have never interpreted content through appearance alone.

They consider where something appears, who presents it, what purpose it seems to serve, what they know about its history and how it relates to other information.

A perfectly realistic fictional image can still be understood as fiction when its context makes that relationship clear. A technically imperfect recording can still carry strong meaning when its source and circumstances are well understood.

Increasing realism therefore does not eliminate context.

It can make context more important because appearance contributes less information about production.

The durable change is not necessarily that people become unable to distinguish synthetic from recorded content visually. It is that **visual or auditory distinction may no longer be sufficient as the primary basis for understanding origin.**

Genuine content remains part of the changed environment

The changing environment also affects material that AI did not generate.

Chapter 9 examined how genuine content can itself become easier to question when synthetic alternatives are known to exist. This consequence may remain important even if technologies for describing or tracing generated content improve.

Once the possibility of convincing synthesis becomes part of public understanding, people know that apparent evidence can have more than one possible production history.

That knowledge cannot simply be removed.

Yet it does not follow that all content becomes permanently doubtful.

Context, provenance, multiple sources, direct experience and established relationships can continue to support ordinary understanding. People already navigate many situations in which complete certainty is unavailable.

The challenge is therefore not to restore an earlier world in which content supposedly spoke entirely for itself. Such a world never fully existed.

It is to understand the additional uncertainties that synthetic production introduces into an environment that was already mediated, selective and imperfect.

Shared understanding can also evolve

Chapter 11 considered the implications of synthetic content for shared reality.

Those implications should not be understood as a prediction that common informational ground must disappear.

Social conventions can adapt. Institutions can develop new practices. Audiences can become more familiar with generated material. Technologies for communicating provenance may change. New expectations may emerge around particular forms of content.

At the same time, disagreement will remain.

People will continue to interpret information differently. Some will have greater access to contextual information than others. Media environments will differ across countries and communities. Technological familiarity will remain uneven.

Shared reality has never required eliminating those differences.

Its continuing challenge is to maintain enough common reference points for people to distinguish disagreement about meaning from disagreement about whether the underlying object or event exists in the form being discussed.

Synthetic content adds complexity to that challenge without determining its outcome.

Human judgement remains necessary because context cannot be automated away

Throughout this Guide, human judgement has remained central.

That does not mean humans will always correctly identify generated content. Nor does it imply that individuals should personally resolve technical questions that may require specialist knowledge.

Human judgement matters for a different reason.

Content always appears within circumstances.

Its significance depends on what it claims to represent, where it appears, what purpose it serves, what consequences attach to it and what remains unknown about it. These are contextual relationships, not merely technical characteristics.

Technology can provide information about production. Institutions can establish practices. Systems may help describe provenance. But the meaning attached to those signals remains connected with human purposes and situations.

The same generated image can be trivial in one context and highly consequential in another.

The same uncertainty about origin can be irrelevant to one interpretation and central to another.

No single technological classification can capture all of those differences.

The enduring issue is not AI alone

It is possible to describe the rise of synthetic content as a story about increasingly capable artificial intelligence.

That description is incomplete.

The deeper story concerns the human information environment.

People create, communicate, represent, interpret, share and remember through content. Technologies participate in those processes, but their significance arises through how they alter relationships among people and information.

Generative AI changes what can be produced.

It changes who can produce some forms of material and how quickly.

It changes the range of possible relationships between what content looks like and how it came into existence.

It can increase the amount and variation of material in circulation.

It can complicate familiar assumptions about authorship, authenticity and recorded presence.

And, because people know these capabilities exist, it can change how even genuine material is received.

These are not simply properties of AI systems.

They are characteristics of the environment humans increasingly inhabit.

Living with the change

The title of this Guide deliberately uses the phrase *Living with AI-Generated Content*.

Living with something is different from treating it as a temporary novelty.

It suggests an environment in which synthetic material becomes one part of ordinary informational life alongside human-created, recorded, edited, simulated and mixed content.

That environment will continue to change.

Specific generation techniques will improve or disappear. New media forms may emerge. Social conventions will develop. Expectations about production may shift. The boundaries that seem important today may be redrawn.

For that reason, the most durable understanding does not depend on memorising the temporary characteristics of synthetic content.

It depends on retaining the conceptual distinctions that remain meaningful as technologies evolve.

What does the content appear to represent?

What is known about where it came from?

What relationship exists between its appearance and its production?

What does authenticity mean in this particular context?

What remains uncertain?

These are not instructions for reaching a prescribed conclusion. They identify the dimensions along which the information environment itself has become more complex.

The central change can therefore be stated simply.

For much of everyday life, people could often move from seeing, hearing or reading content to an intuitive assumption about how it had been produced.

That assumption was never infallible.

Increasingly, it is no longer sufficient.

Text may be written, assisted or generated. Images may be photographed, modified, synthesised or combined. Voices may be recorded, transformed or created. Video may document, reconstruct, simulate or blend several forms of production.

The surface remains meaningful, but it no longer necessarily reveals the history beneath it.

Living with AI-generated content means recognising that difference while continuing to interpret, communicate and understand within the human world in which the content appears.

AI may generate the material.

Meaning remains a human question.

Key Ideas from This Guide

AI-generated content is not simply a new category of media.

It contributes to a wider change in the information environment.

Several ideas run through the Guide.

Appearance and origin are increasingly separable.

A photograph-like image need not originate with a camera. A recognisable voice need not correspond to an act of speaking. Fluent text does not disclose how it was composed.

Human-created and machine-generated are not always opposing categories.

Content can pass through multiple stages of human and machine contribution. Assistance, modification, generation, selection and editing may coexist within one production process.

The production process can be a continuum.

Knowing that AI was involved may not reveal how much it contributed, where it contributed or how significant that contribution was.

Different media carry different historical expectations.

Text, photographs, voices and video have not traditionally communicated origin in the same way. Generative AI therefore changes the informational significance of each differently.

Authenticity is not a synonym for truth.

Human-created content can be false. Generated content can communicate accurate information. Authenticity may concern source, authorship, representation, intention or production.

Provenance concerns history, not merely appearance.

Understanding where content came from and what happened to it during production and circulation can reveal information that the finished surface does not.

Representation and occurrence can be separated.

A realistic depiction of a person, place or event does not necessarily establish that the circumstances shown physically occurred in the form presented.

Scale matters.

Generative technologies can affect the information environment not only through what one item can contain but through how rapidly and extensively variations can be produced and circulated.

Information, expression and simulation can overlap.

Similar technologies can be used to document, explain, illustrate, persuade, entertain or simulate. Purpose and context therefore remain important.

Synthetic content can affect how genuine content is received.

Once convincing generation is widely understood to be possible, genuine photographs, recordings and other material can themselves become easier to question.

Repetition does not necessarily mean independent confirmation.

Familiarity and apparent multiplicity can influence how content is perceived even when repeated material ultimately has a common origin.

Context changes consequence.

The same technological process can carry very different significance in journalism, entertainment, education, advertising, creative work, workplace communication or personal communication.

Uncertainty about origin may sometimes remain unresolved.

Not every encounter requires a technical determination of how content was produced. The significance of uncertainty depends upon context, purpose and consequence.

Human judgement remains contextual.

Technology can provide information about production and provenance, but the significance attached to those signals remains connected to human purposes and circumstances.

Finally, the information environment itself will continue to change.

Specific generation techniques will improve or disappear. New media forms may emerge. Social conventions will develop. Expectations surrounding authorship and production may shift.

The most durable understanding therefore does not depend upon memorising today's supposed signs of synthetic content.

It depends upon retaining the conceptual distinctions that remain meaningful as technology evolves.

The surface of content still matters.

Increasingly, however, the surface does not necessarily reveal the history beneath it.

AI Guidance — Continuing the Exploration

AIGuide 10 examines one particular dimension of the relationship between people and artificial intelligence: **living within an information environment containing increasing amounts of generated, modified and mixed-origin content.**

Several neighbouring questions are deliberately treated elsewhere within the AISF AI Guidance series.

AIGuide 7 — Knowing When to Trust AI addresses the broader question of trust in AI and its outputs.

AIGuide 10 encounters trust where familiarity, repetition, context and synthetic content affect perception, but it does not attempt to provide the wider architecture for deciding when AI should be trusted.

AIGuide 8 — AI and the Human Decision examines the broader territory of AI and human decision-making.

AIGuide 10 remains more narrowly concerned with the interpretive conditions surrounding synthetic content before or alongside later decisions.

AIGuide 9 — AI, Truth and Verification examines detailed factual verification and factual correctness.

AIGuide 10 distinguishes those questions from origin, authenticity, provenance and interpretation.

These are neighbouring territories rather than dependencies.

AIGuide 10 is intended to stand independently. Readers who wish to explore trust, human decision-making or factual verification more deeply can continue into the corresponding Guides.

All AISF AIGuides are provided to support informed curiosity, clearer understanding and independent thinking.

Future Edition Updates & User Submissions

AI-generated content is developing rapidly.

Generation technologies may become more capable. New combinations of recorded and synthetic material may emerge. Provenance mechanisms may develop. Social and professional

conventions may change. Different jurisdictions may introduce or revise relevant legal and regulatory approaches.

For that reason, AISF treats AI Guidance as an evolving educational resource rather than a permanently fixed description of a changing environment.

Future editions may incorporate material developments where they meaningfully affect the concepts examined in this Guide.

AISF also welcomes observations from readers concerning developments that may warrant future consideration, areas where additional explanation could improve clarity, emerging questions surrounding AI-generated content and potential subjects for 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 the development process, including determination of purpose, scope, architecture, editorial standards, substantive acceptance and final publication.

AISF also acknowledges the wider community of researchers, educators, practitioners, creative professionals, institutions, developers and users whose continuing engagement with artificial intelligence contributes to broader understanding of generated content and the changing information environment.

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 should not be interpreted as professional, legal, financial, medical, technical, regulatory or other specialised advice.

The Guide does not recommend or endorse any particular artificial intelligence system, provider, product, platform, generation technology, detection technology, provenance mechanism, organisation or course of action.

Artificial intelligence technologies, services and capabilities can change rapidly. Their operation may differ between providers, models, versions, jurisdictions and settings. Information or observations contained in this Guide should therefore not be interpreted as describing every AI system, generated-content technology or circumstance.

Examples are illustrative and are intended to assist understanding of broader concepts.

Discussion of authenticity, provenance, origin, synthetic content, generation, modification and related concepts is educational and conceptual. Terminology, legal classifications, platform policies, professional standards and regulatory requirements may differ between jurisdictions and may change over time.

This Guide does not provide a technical method for determining whether particular content is AI-generated, altered, authentic or genuine.

Where questions concerning content origin, authenticity or factual accuracy may contribute to decisions carrying material consequences, readers should consider the relevant circumstances and, where appropriate, consult current authoritative sources or suitably qualified professionals.

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: Living with AI-Generated Content

Subtitle: Understanding Origin, Authenticity and Interpretation in an Increasingly Synthetic Information Environment

Series: AI Guidance

Guide Number: AIGuide 10

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, generated and modified content, provenance practices and the wider information environment.

Back Cover

Living with AI-Generated Content

A photograph may look real without having been photographed.

A familiar voice may be heard saying words the person never physically spoke.

A piece of writing may contain contributions from both a human and AI.

A video may combine recorded and generated elements so seamlessly that its production history is not visible from the finished result.

What happens when content no longer tells us how it came into existence?

AIGuide 10 — Living with AI-Generated Content examines the changing relationship between content, origin, authenticity and human interpretation.

Across fourteen chapters, the Guide explores the continuum between human and machine creation, text, images, voices and video, authenticity, provenance, synthetic representations, generation at scale, shared reality, context and interpretation under uncertainty.

It does not assume that AI-generated content is inherently deceptive.

Nor does it provide a technical formula for detecting it.

Instead, it examines a more enduring question:

How do people understand content when its surface increasingly reveals less about the history beneath it?

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