

Using Generative AI Responsibly In Electronic Disclosure

This guide explores the responsible application of generative AI in electronic disclosure, covering prompt engineering, human verification, and defensibility. It addresses the guidance landscape, common mistakes, and practical considerations for UK lawyers. Topics include relevance, privilege, summarisation, and redaction support.

e-discovery.uk, Computer Forensics Lab. Published 2026-08-30.

<https://e-discovery.uk/library/using-generative-ai-responsibly-in-electronic-disclosure>

GENERATIVEAIINDISCLOSURE · A GUIDE FOR UK LAWYERS Using Generative AI Responsibly in Electronic Disclosure Prompted Review, Human Verification and Defensibility Under PD 57AD COMPUTER FORENSICS LAB

§ ABOUT THE AUTHOR PREPARED BY COMPUTER FORENSICS LAB E-DISCOVERY TEAM ESTABLISHED 2007 · LONDON ISO 17025-ALIGNED PROCEDURES GENAI-ASSISTED REVIEW WORKFLOW S

§ CONTENTS In this guide 01 Executive summary 02 The problem in plain English: a reader that never tires and some time s invents 03 The guidance landscape: ILTA, the Law Society, the judiciary and Ayinde 04 Where GenAI fits: relevance, privilege, summarisation and redaction support 05 Prompt engineering as methodology: drafting, testing and version in g 06 Verification: sampling AI outputs against human judgement 07 Transparency, data protection and the protocol 08 Source architecture: where else the evidence lives 09 Worked examples 10 Common mistakes and technical limitations 11 Questions to ask · Suggested wording 12 Checklist and red flags · When to involve a digital forensic expert 13 Frequently asked questions 14 Glossary · References · Disclaimer · How a specialist laboratory can assist

§ 01 · ORIENTATION Executive summary THE HEADLINE POINT: GENAI READS AND CHARACTERISES AT SCALE UNDER INSTRUCTIONS YOU WRITE: THE INSTRUCTIONS ARE TESTED LIKE SEARCH TERMS, THE OUTPUTS VERIFIED LIKE A JUNIOR 'S' WORK, AND THE LAWYER REMAINS ACCOUNTABLE FOR ALL OF IT

§ 02 · FIRST PRINCIPLES The problem in plain English: a reader that never tires and some time s invents

§ 03 · THE RULES OF THE ROAD The guidance landscape: ILTA, the Law Society, the judiciary and Ayinde

§ 04 · THE USE CASES Where GenAI fits: relevance, privilege, summarisation and redaction support

§ 05 · INSTRUCTIONS AS METHOD Prompt engineering as methodology: drafting, testing and version in g

§ 06 · CHECKING THE READER Verification: sampling AI outputs against human judgement

§ 07 · THE PAPER WORK AND THE DATA Transparency, data protection and the protocol

§ 08 · THE WIDER MAP Source architecture: where else the evidence lives QUESTION
DECISION QUALITY PROMPT VERIFICATION PROCESSING SCOPE FILE RECORDS
INVENTORY LAYERS HUMAN TEXT- LOG LAYER

§ 09 · IN THE WILD Worked examples EXAMPLE1 ·
THE PROMPTED FIRST PASSTHAT SURVIVED ITS AUDIT EXAMPLE2 ·
THE SUMMARY THAT INVENTED A CONCESSION EXAMPLE3 ·
THE PRIVILEGE FLAG THAT ACCELERATED, AND NEVER DECIDED

§ 10 · WHERE IT GOES WRONG Common mistakes and technical limitations Common mistakes
Technical limitations

§ 11 · INTERROGATORIES & DRAFTING AIDS Questions to ask · Suggested wording Ask
your client Ask your opponent Ask your e Discovery / forensic provider SUGGESTED
WORDING · GENAI PARAGRAPH FOR THE DRD / PROTOCOL

§ 12 · QUICK CONTROL Checklist and red flags · When to involve a digital forensic expert The
GenAI checklist Red flags When to involve a digital forensic expert

§ 13 · COMMON QUESTIONS Frequently asked questions Is it actually permissible to use
generative AI in court-ordered disclosure? How is this different from the TAR we already
use? What is "hallucination" and how do we protect against it in disclosure work? Do we have to
tell the other side we used AI? Can AI make privilege calls? Our client's data is highly
sensitive. Does that rule GenAI out?

§ 14 · REFERENCE Glossary Sources and authoritative references DISCLAIMER

§ HOW A SPECIALIST LABORATORY CAN ASSIST Working with Computer Forensics Lab
Speak to a forensic examiner, not a salesperson. INSTRUCT THE LAB
NEW ENQUIRIES EMAIL - DISCOVERY

Questions and answers

Is it actually permissible to use generative AI in court-ordered disclosure?

Yes: PD 57AD's technology duty accommodates it, the ILTA best-practice guide (September 2025) exists precisely to structure it, and the Law Society's practical guidance points practitioners to that framework: the conditions are methodology, verification, documentation and disclosure of use, not permission. What the authorities punish: Ayinde's line: is unverified reliance, not the tool.

How is this different from the TAR we already use?

TAR ranks by learned resemblance to coded examples; GenAI reads against written instructions and returns assessments with reasons: instructions you can draft, test and version like search terms: guide 89's constitution carries over, with the prompt file joining the training record. In practice they combine: ranking for triage, prompted reading for assessment, humans for decisions.

What is "hallucination" and how do we protect against it in disclosure work?

Confident output not grounded in the document: Example 2's invented concession is the disclosure-native form. Protections are structural: citation-required prompts, verification sampling at consequence-scaled rates, and the standing rule that summaries are leads: every load-bearing assertion checked against the source before it enters work product.

Do we have to tell the other side we used AI?

Under PD 57AD's cooperation model, effectively yes for outgoing disclosure: methodology transparency is the expectation the ILTA guide operationalises, and the DRD is the venue: what is disclosed is workflow and verification, not your prompts' work product or document-level outputs. Silent use discovered later is the most avoidable dispute in this guide.

Can AI make privilege calls?

It can flag, queue and alert: Example 3's architecture: and the determination remains a lawyer's on a read document: the line every current authority draws and the one place verification is never economised. An AI- accelerated privilege review is good practice; an AI-adjudicated one is a waiver argument waiting to happen.

Our client's data is highly sensitive. Does that rule GenAI out?

It rules the posture in: platform terms, residency, security and process or obligations assessed against UK GDPR and confidential it y duties before anything is processed: with deployment models ranging from restricted enterprise environments to not-at-all: §7's analysis decides, documented either way. Sensitivity shapes the how; it does not by itself answer the whether. cflab. u k · e-disc ove r y. u k ©2026 Computer Forensics Lab Ltd ·cflab.uk ·e-discovery.uk ·info@cflab.uk ·+44 (0)20 7164 6915 Page 15 of 17

Source: <https://e-discovery.uk/library/using-generative-ai-responsibly-in-electronic-disclosure>. Free to read and share with attribution to Computer Forensics Lab, e-discovery.uk.