



AI Content and Knowledge Work Guidelines

Working standards for using AI to write, research, analyze and create without losing accuracy, voice or trust.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 03 / EMBER & LILAC

EXECUTIVE WORKING EDITION - 08 OCT 2026

03

AT A GLANCE

Working standards for using AI to write, research, analyze and create without losing accuracy, voice or trust.



Knowledge workers now use AI to draft, summarize, research, analyze and design. Done well, it saves hours. Done carelessly, it ships invented facts, leaks confidential material and erodes the voice and credibility of the organization. These guidelines set practical standards for AI-assisted content: where it helps, how to prompt and verify, how to handle sources and intellectual property, when to disclose, and how editors and managers review the work. They include checklists, a review log and a quick reference. They are general guidance, not legal advice.

WHO IT IS FOR

Marketing, communications, research, consulting, legal, finance and operations teams, their managers and editors, and anyone producing documents, analysis or media for colleagues, clients or the public.

TIME TO USE

60 minutes to read; one team session to agree your review and disclosure rules



What you will have at the end

- A clear map of where AI helps, where it needs care and where it should not be used for content
- A verification routine for facts, figures, quotes and citations
- Rules for sources, intellectual property, confidentiality and disclosure
- An editorial review path scaled to the audience and the stakes
- A shared voice standard, review log and quick reference for the whole team



Contents

01	What these guidelines cover and why content needs its own rules	03
02	Six principles for AI-assisted knowledge work	04
03	Approved, approval-required and prohibited content uses	05
04	How to prompt, check and cite	07
05	Confidentiality, intellectual property and attribution	08
06	When to say AI helped, and how to keep a human voice	10
07	Editorial review by stakes, roles and handling errors	11
08	Training, review cadence, checklists and the one-page reference	13
09	A first quarter that produces evidence	14
10	Terms used in this publication	15



PURPOSE AND SCOPE

What these guidelines cover and why content needs its own rules

01

II

Content errors travel under your name, so the standard must be set by the quality you promise, not by the speed the tool allows.

Content is where AI is most tempting and where mistakes are most visible. A single invented statistic in a client report, a fake citation in a legal memo or a leaked client name in a prompt can cost more than a year of saved time. These guidelines assume AI is useful and focus on keeping the human professional responsible. They sit beside your general responsible use policy and your data handling rules, and where they conflict, the stricter rule applies.

In scope

- Written work: articles, emails, proposals, reports, presentations, social posts, scripts and web copy.
- Research and analysis: literature summaries, market scans, data analysis, code for analysis and financial or legal drafts.
- Media: images, audio, video, design and translation.
- Knowledge assets: playbooks, training material, FAQs, internal wikis and help content.
- Anyone producing these items for colleagues, clients, regulators or the public, whether employee or contractor.

NOT LEGAL ADVICE

Copyright, trademark, privacy, advertising, professional conduct and AI transparency rules differ by country and sector, and several are changing. Have counsel confirm how they apply to your content before you publish the adapted version of these guidelines.

WORKING QUESTION

What is the worst factual error that could leave your team under your organization's name, and what stops it today?

PRINCIPLES

Six principles for AI-assisted knowledge work

02

Use AI to speed up thinking and drafting, never to replace the judgment and accountability that clients pay for.

PRINCIPLE	IN PRACTICE
The author is accountable	The person who publishes or sends the work owns every claim in it, however it was produced.
Facts are verified	Figures, names, dates, quotes, legal references and citations are checked against primary sources.
Confidence is not evidence	Fluent text is not proof. Treat every unsourced claim as a lead to confirm.
Voice is ours	Content should sound like the organization and the author, not like generic machine prose.
Sources and rights are respected	Do not copy protected material, and credit and license properly.
Honesty about the process	Disclose AI use when the reader would reasonably want to know.

These principles lead to two working habits. First, treat AI output as a draft from a fast, confident junior colleague who has read widely but cannot be trusted with facts and who knows nothing about your client unless you tell them. Second, keep your own thinking in the loop: decide the argument, the audience and the point of view first, then use the tool to explore, structure and polish, not to originate every idea.

COMMON MISTAKE

Asking the tool to write the whole piece from a single line, then editing lightly. The result is generic and error-prone. Better: decide the outline and key points yourself, ask for a draft of one section at a time, and rewrite in your own voice.

WORKING QUESTION

In your team's last AI-assisted deliverable, which parts were the human's thinking and which were the tool's?

WHERE AI FITS

Approved, approval-required and prohibited content uses

03

Match the tool to the task and the stakes: generous for drafting, strict for anything someone will rely on.

LANE	USES	CONDITIONS
Approved	Outlining, brainstorming, rephrasing, tone adjustment, proofreading, summarizing your own notes	Approved tool; no restricted data; author reviews all output
Approved	First drafts of internal or low-stakes material	Author edits, verifies facts and owns the result
Approved	Structuring research questions and suggesting where to look	Findings come from primary sources, not from the tool's recall
Approval required	Client deliverables, published thought leadership, advertising and regulated communications	Second-person review; disclosure rules apply; source list kept
Approval required	Use of client or personal data, or any confidential source material	Approved tool for that data tier; client contract permits it

LANE	USES	CONDITIONS
Approval required	Synthetic images, voice, video or likenesses of real people	Consent, rights and labeling check
Not allowed	Fabricating quotes, testimonials, reviews, case studies or data	No exceptions; hypothetical examples must be labeled as such
Not allowed	Presenting AI-produced legal, medical, financial or safety advice without qualified review	Professional must review and adopt

Different teams will draw the lines differently. A legal team may restrict AI to non-substantive editing, while a marketing team may allow broader drafting. Record these team variations in one shared document so that staff moving between teams are not surprised. The aim is predictable behavior, not uniformity for its own sake.

Analysis deserves particular care. When AI is used to analyze a dataset, write a formula or produce code for a calculation, the output looks authoritative because it is precise. Check the method as well as the answer: confirm the tool used the right columns, filters and time period, rerun the calculation independently on a sample, and keep the working file. A confidently wrong chart in a board paper is harder to retract than a clumsy sentence, because people remember the picture.

WORKING QUESTION

Does your team have a written list of content that must never be AI-drafted, and does everyone know it?

VERIFICATION

How to prompt, check and cite

04



Better prompts raise the floor, but only verification protects the ceiling of your reputation.

Prompting habits that help

- State the audience, purpose, format and length, and say what a good result looks like.
- Supply the source material you want used, and instruct the tool to rely only on it and say when it cannot find an answer.
- Ask for the reasoning or evidence behind claims, then check that evidence yourself.
- Work in short iterations: outline, section, revision, rather than one giant request.
- Remove client names, personal data and other detail the task does not need before pasting anything.



Models can state invented facts with complete confidence, including plausible citations to papers, cases or reports that do not exist. This is called hallucination, and it appears in summaries, research, translation and code. Do not ask the same tool to confirm its own claims and treat agreement as proof. Verification must go back to the original document, dataset, filing, statute or expert.

**01 Mark every claim**

Highlight figures, names, dates, quotes, legal or regulatory references, product claims and any statement of cause.

02 Trace each to a source

Find the primary source and confirm that it says what the draft says. Replace or remove anything you cannot trace.

03 Check the numbers

Recalculate totals, percentages and comparisons. Confirm units, dates and the basis of comparison.



04 Read for omissions and tone

Ask what is missing, what is overstated and whether the claim would survive a skeptical expert.

05 Record the check

Initial the review log with date, sources used and reviewer, so the work can be defended later.

WORKING QUESTION

For the last number you published, can you open the primary source in under two minutes?

SOURCES AND RIGHTS

Confidentiality, intellectual property and attribution

05

What goes into a prompt and what comes out of the tool both carry legal and contractual obligations.

Start with confidentiality. Client files, unreleased results, personal data and privileged material are subject to contract and law. Use only tools cleared for that class of information, and check your client agreements, because some forbid sharing material with third-party AI services or require consent. If in doubt, ask before pasting. The information-tier rules in your responsible use guidelines apply to every prompt.

Intellectual property is less settled. Rules on the copyright status of AI-generated material, on training data and on reproduction of protected works vary by jurisdiction and are the subject of ongoing litigation and legislation. Treat output as potentially unprotected, avoid asking tools to imitate a named living creator or reproduce protected text, and review vendor terms on ownership and indemnity before using output commercially at scale. Ask counsel for current guidance.



SITUATION	WORKING RULE
Quoting another source	Quote exactly from the original, cite it, and stay within fair-use or licensed limits
Using generated images or music	Check the tool's license for commercial use and avoid recognizable people, brands or styles
Summarizing a paywalled or licensed report	Check the license allows use in AI tools and that the summary does not substitute for the original
Including client material in a deliverable	Use only with permission and keep confidential details out of third-party tools
Reusing prior team work as examples for the tool	Remove client identifiers and confirm you have the right to reuse it

ATtribution NORM

Cite the real source of every fact, quote or idea you did not originate. Never cite the AI tool as a source of facts. If the tool pointed you to a source, read the source and cite it, not the tool.

WORKING QUESTION

Which client agreement in your portfolio restricts use of AI tools, and does your team know it?

DISCLOSURE AND VOICE

When to say AI helped, and how to keep a human voice

06

Readers forgive assistance more readily than concealment, and they notice generic text immediately.

Disclosure should follow the reader's reasonable expectations. A spell-checked email needs no label. A research report whose first draft was machine-written may warrant a line in the method note. Synthetic images of real-looking people, cloned voices and automated customer conversations nearly always need clear labeling, and some jurisdictions or platforms require it. Contracts, publication standards and professional rules may also set the terms. When in doubt, disclose.

CONTENT	DISCLOSURE APPROACH
Internal drafts and notes	No label required; the author is accountable
Client deliverables	Follow the contract; note AI assistance in the method where relevant or requested
Published articles and marketing	Follow platform and legal rules; add a short note when AI contributed substantially
Synthetic media of people or realistic scenes	Label clearly and get consent where a real person is depicted
Chatbots and automated messages	Tell the person they are dealing with AI and offer a route to a human

Protecting voice and quality

- Keep a short style guide with tone, vocabulary, banned phrases and examples of work you consider excellent.
- Edit for specifics: real examples, concrete numbers from your own data, named decisions and clear positions.
- Cut filler openings, repetitive structures, vague claims and inflated adjectives that signal machine prose.
- Read aloud. If it does not sound like the author, rewrite it.

- Do not publish at volume simply because it is cheap. Every item still costs reader attention and your credibility.

Translation and localization need a specific rule. Machine translation is useful for understanding a document, but a published translation carries your name in a language you may not read. For anything external, ask a fluent reviewer to check meaning, tone, legal terms and cultural fit. The same applies to summaries of long documents: a summary can drop the one clause that matters, so reviewers who rely on a summary should know what was left out and who checked it against the full text.

WORKING QUESTION

Would a long-standing client recognize your team's writing as yours with the logo removed?

REVIEW AND ESCALATION

Editorial review by stakes, roles and handling errors

07

Scale review to the audience and the consequence, and make correcting an error cheap and expected.

STAKES	EXAMPLES	REVIEW REQUIRED
Low	Internal notes, brainstorm, drafts	Author verifies and owns for the author
Medium	Client emails, reports, proposals, posts	Author verification plus peer or editor review before release
High	Regulated communications, legal and financial advice, public statements, safety information	Subject expert and approver sign-off; review log kept; counsel as needed
Critical	Content that could mislead many people or cause significant harm and a correction plan wrong	Executive sign-off, staged release

Roles

- Author: drafts, verifies and takes responsibility. Cannot delegate accountability to the tool.
- Reviewer or editor: checks accuracy, tone, rights and disclosure, with authority to stop release.
- Team lead: sets the review level for the team's content types and trains new staff.
- Guideline owner: maintains this document, the approved tool list and the review log template.
- Everyone: reports errors and near misses without waiting to be asked.

01 Correct quickly

If you find an error in published content, tell the owner, fix it and add a visible correction when readers relied on it.

02 Contain

Check whether the same tool, prompt or source produced other errors, and pause that workflow.

03 Record

Write a short incident note: what, where, how found, who is affected, cause, fix.

04 Escalate where needed

Involve counsel and communications for client impact, data exposure, defamation, rights claims or regulatory content.

05 Improve

Update prompts, checklists or training, and share the lesson without naming the author.

WORKING QUESTION

When was the last correction your team made to published work, and how quickly was it spotted?

TRAINING AND QUICK REFERENCE

Training, review cadence, checklists and the one-page reference

08



Skills and habits are what make these guidelines real, so train them, measure them and review them on a schedule.

Run a 60-minute workshop for every team that produces content, using that team's own tasks. Cover prompting, verification, data tiers, rights and disclosure, then have participants fix a deliberately flawed AI draft with errors you planted. Repeat annually, and brief new starters in their first month. Ask each person to acknowledge the guidelines in writing, and keep the record. Review the document every six months in the first year, and again after any major tool change, new regulation or incident.



Measure whether the guidelines are working with a few simple signals rather than a large dashboard. Sample a handful of published items each month and count errors found after release. Track how often reviewers send work back, since a rate of zero suggests nobody is looking. Ask clients and colleagues whether quality and turnaround have improved. Time saved matters only if accuracy holds, so report both together to the sponsor and adjust the review levels accordingly.



Before you publish AI-assisted work

- Every figure, name, date, quote and citation traced to a primary source
- No confidential or personal information placed in an unapproved tool
- Text reads in our voice and carries specific examples, not generic claims
- Rights checked for images, music, quoted material and client content
- Disclosure applied where the reader, contract or law expects it
- Reviewer sign-off recorded at the level the stakes require
- Author is prepared to defend every claim



Review log fields

- Title, author, reviewer, date and audience
- Tool used, data tier of inputs and the stakes level
- Claims checked and sources used



- Changes made and issues found
- Disclosure decision and approver

QUICK REFERENCE	RULE
Draft	Use AI freely for structure and first drafts with approved tools
Verify	Trace every claim to a primary source; never trust the tool to check itself
Protect	No confidential or restricted data in unapproved tools; minimize detail
Disclose	When the reader or rules expect it; always for synthetic people and bots
Own	You are accountable for what you publish, in your voice, with a reviewer for higher stakes

WORKING QUESTION

What single check, if made routine tomorrow, would remove the most risk from your team's published work?

30 / 60 / 90

A first quarter that produces evidence

09

Introduce the standards by fixing the highest-risk habits first, then building shared voice and a review rhythm.

First 30 days: ground it

- Adopt the approved, approval-required and not-allowed lists and adapt them by team
- Publish the approved tool list and the data tier rules
- Brief managers and editors on the verification routine
- Distribute the quick reference and pre-publication checklist

Days 31 to 60: test it

- Run the 60-minute workshop with each content team
- Start the review log for client deliverables and public content
- Write the team style guide with examples of good work
- Collect written acknowledgement from everyone in scope



Days 61 to 90: decide

- Audit a sample of published items against the log and checklist, and share the findings openly with the team
- Review incidents and corrections, and update prompts and rules
- Check client contracts for AI restrictions and align practice
- Set the six-month review date and annual training schedule



GLOSSARY

Terms used in this publication

10

TERM	MEANING
Hallucination	A fluent output that is false or unsupported, such as an invented statistic, quote or citation.
Prompt	The instruction and context given to an AI tool to produce an output.
Primary source	The original document, dataset, filing or statement, as opposed to a summary or secondhand account.
Grounding	Instructing a tool to base its answer on supplied source material rather than on its general training.
Verification	Checking a claim against a primary source or qualified expert before relying on it.
Disclosure	A statement that AI contributed to a piece of work or that a person is interacting with AI.
Synthetic media	Images, audio or video generated or altered by AI, including realistic depictions of people.



TERM	MEANING
Data tier	A sensitivity level for information that determines which tools may receive it.
Style guide	A short document defining tone, vocabulary and standards, with examples of good work.
Review log	A record of who reviewed a piece of work, what was checked and what was decided.
Fair use	A legal doctrine in some jurisdictions allowing limited use of protected material; scope varies and needs counsel.
Automation bias	The tendency to accept machine output without enough checking.

A note on use

These guidelines are general operating guidance and not legal advice. Confirm current copyright, privacy, advertising, professional and AI-specific requirements with qualified counsel.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.