

The Altlaw Guide to Prompt Engineering

Making aiR for Review work for you.



When it comes to the adoption of new tools and technologies, the legal industry is, as a rule, rather slow on the uptake. With GenAI, this has not been the case.

Let's explore what makes this new wave of Generative AI tools so intriguing.

Contents

- 03** [On the adoption of AI](#)
- 05** [What is aiR for Review?](#)
- 07** [The importance of your Review Protocol](#)
- 09** [How to write a Review Protocol](#)
- 12** [Prompt engineering for aiR for Review](#)
- 17** [Prompt writing best practices](#)





On the adoption of AI

Relativity's **aiR for Review** - the latest in generative AI tools to hit the eDiscovery market, has been met with widespread intrigue and dare we say, quiet excitement!

This is a far cry from the reception granted to the emergence of Active Learning, some 10 years ago.

So, what has caused this sea change in attitudes towards AI within the legal industry?

The first factor we must take into account is trust.

When Active Learning was first released, there were no other tools on the market with which to compare. No understanding of how these tools might affect our processes, results and our bottom line.

Now we have an established trust with Active Learning technologies and plenty to choose from. This helps us to be more open-minded when it comes to adopting new tools.

“The more care and consideration you put into your prompt, the more successful your review will be, ergo the more time and money you will save.”

The second factor is familiarity.

Active Learning was an unseen technology when it was first released, but the same cannot be said for generative AI. Chat-GPT has taken the world by storm in the past months, and as such, people are already very familiar with the technology and how it works.

This results in tools based on this technology being less frightening as people understand them on a deeper level and are more comfortable using them.

Read on to find out more about aiR for Review, how to write a Review Protocol and how to engineer that document into a prompt that will boost your efficiencies and save you time and money.



What is aiR for Review?

aiR for Review is Relativity's latest AI tool. Based on OpenAI's GPT-4 technology it empowers you to review with greater speed and efficiency, helping you win faster with AI.

aiR for Review uses large language modelling algorithms to predict which of your documents is likely to be relevant to your case based on the relevance criteria and case information you provide the model.

Once the model has been trained (you have entered your prompt) it reviews your documents and predicts which are likely to be relevant/hot and provides reasoning and citations to support its predictions.

By grounding the model in citations, we help mitigate the risk of hallucinations. This is also aided by the fact that the system is entirely closed, so no data other than the information you provide can enter or leave your workspace. This also ensures your data remains secure.

There is a specific order that the aiR tool completes its prediction process to maximise the information you gain and also the model's likelihood of producing a correct prediction.

This order – known as 'Bottom-up' reasoning alongside the algorithm's 'chain-of-thought' processing means that the algorithm collates evidence and uses all of the evidence it has created and produced to make its prediction. It looks a little something like this...



Here we see that the model first looks for citations within the text of a document that may match any of the information provided in your Review Prompt.

Once the model has located a citation it is externally verified by the model and then rationale, explaining why the model believes this citation to be relevant is created.

Next the model is asked to think of any reason why this document may not be relevant, even though a citation has been found.

Finally, based on the citation, rationale and considerations the model will provide an overall prediction for relevance, giving the document a score between 1 and 4. 1 is Not Relevant and 4 is Highly Relevant.

Manual Review

- Time Consuming
- Labour Intensive
- Full Review Required
- Risk of Human Error
- Tried and Trusted

Active Learning x aiR for Review

- Time Sensitive
Prioritised Review
- Instant Training
- Labour Saving
- Can Handle Huge
Quantities of Data
- Cutting-edge
Technology

[Read our aiR for review blog post](#)



The importance of your Review Protocol



In any eDiscovery case, a well-written, thorough review protocol is a must-have if you want to ensure success. The more detailed your protocol, the better understanding your reviewers have of the case and the more accurate they can be.

A well-considered review protocol is also extremely important when using AI, namely aiR for Review, to review documents as this is what trains the AI algorithm that will predict the relevance of your documents.

The more care and consideration you put into your prompt, the more successful your review will be, ergo the more time and money you will save.

Drafting a comprehensive document review protocol provides the best method for ensuring a high-quality and consistent document review, by proactively informing document reviewers of your expectations and delivering relevant information and directions.

Your document review protocol will serve as a directing document, stating objectives, providing case history, reference materials, and explaining coding instructions to guide the document review team.

Although, document reviewers will still have questions after a briefing call from lawyers explaining the engagement, a thoughtfully drafted review protocol will reduce questions and increase efficiency by proactively addressing probable questions through detailed instructions.


“The importance of iterative testing while prompt engineering cannot be over stated.”



How to write a Review Protocol

Although each matter will require tailored sections, typically, a document review protocol should cover:

- Introduction statement
- Document coding instructions
- Case background
- Privileged Documents
- Objective statement
- Document coding reference document examples
- Case context reference documents
- Corporate document services



Introduction Statement:

Begin with a brief introduction to the case and address the scope of the upcoming review.

Case Background:

Provide document reviewers with details about the matter's inception through to where the matter currently stands. Here, document reviewers will be able to learn the case's procedural history, internal developments, case theories, and outside interactions with opposing counsel that may provide additional context to their assignment.

Objective Statement:

Communicate the expectations of the review. The best objective statements provide the document reviewer with a clear purpose for their engagement and, ideally, will be referred to daily

Case Context Reference Documents:

Explain how document reviewers should actually code documents. While each case is different, coding instructions should follow a general format to promote an efficient review. Coding instructions should be drafted to cover the broadest fields to be coded first, working your way to the most specific fields to be coded last.

Privileged Documents:

Good practice is to identify in advance lawyers and/or law firms involved in the matter so that reviewers are alerted to potential privilege issues.

Document Coding Reference Examples:


Include examples of coded documents. Examples of properly coded documents drastically shape document reviewers' impressions on what is a properly coded document, and thus should be selected with significant consideration.

[Read the in-depth guide](#)



Prompt Engineering for aiR for Review

Provide Specific Context



Prompt Testing

Provide Good and Bad Examples



Once you have created your review protocol you have the perfect basis for your aiR for Review prompt.

Unlike Active Learning, which is trained by reviewing documents and learning from reviewer decisions, aiR for Review is trained instantly via the input of this prompt, putting great weight on the information you provide at the beginning of your review process.

What information do I need to provide?

Case Summary:

- Matter Overview – Taken from your review protocol
- Custodians and reference documents
- Noteworthy Entities – Whistleblowers and key actors
- Noteworthy Terms – State and define these
- Additional Context– Anything of note that hasn't been covered by the previous points.

Relevance:

- Relevance Criteria – Include keywords, phrases, legal concepts, parties, entities and specific issues relating to your case among other information.

You should also include information on what a not-relevant document might look like/contain.

Key Documents :

- Key Document Criteria – Include key topics, key people, types of interaction, personal email addresses etc.

Once you have provided your initial prompt to the algorithm, you can then begin the prompt engineering process. This is simply the process of editing your prompt based on the accuracy and responses of the algorithm.

Firstly, it is important to perform some iterative testing to ensure that your prompt is capturing all the correct documents.

As aiR for Review is charged on a per-document basis, you don't want to run the tool over the entirety of your documents only to find that you have missed a key piece of information in the prompt and have to re-run the review.

By taking a sample of 50-100 documents that you know well, you can make sure that your prompt is capturing all the relevant information in a small iterative test. We recommend running a fast-track review over these few documents until all the documents are tagged as you would expect them to be.

Once you have completed this step you can run the tool over the rest of your documents with confidence that all your key documents will be coded accurately.

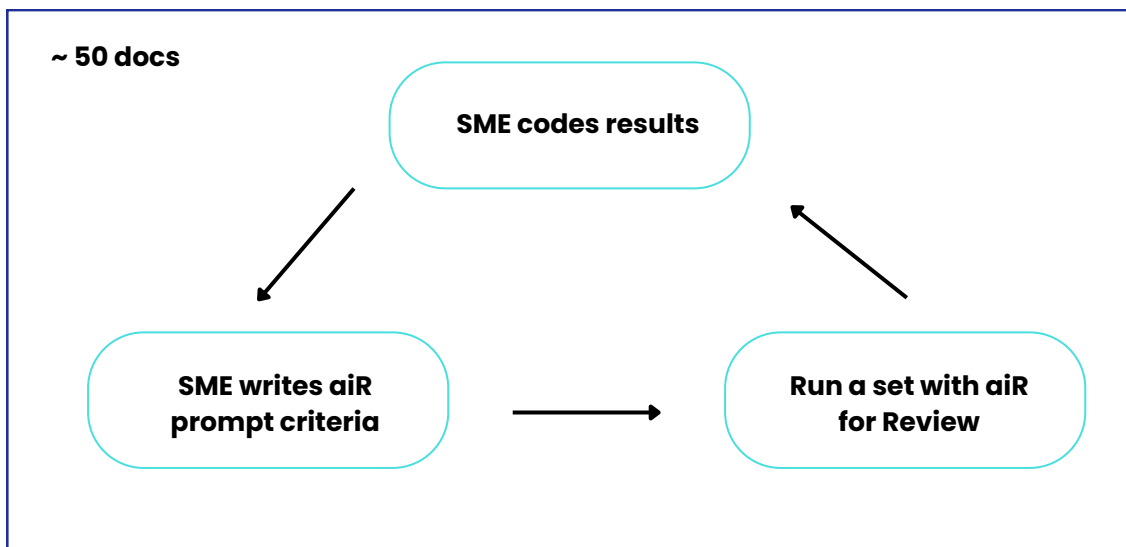
[Chat to a friendly Altlaw expert about how aiR for Review can help you](#)



3 stages of Prompt Engineering:


1. Develop your prompt

This stage should take multiple iterations as you add and subtract relevant information and consolidate your prompt. Here you are testing your prompt criteria on a small group of documents you know well. These should be both relevant and non-relevant.



2. Validate your prompt

At this point, you have a prompt that generates the desired coding results for your well-known documents. Now you want to take a statistical sample of your remaining data to ensure the prompt meets a wider spectrum of issues.



3. Run your review

Once your prompt has been validated and proven successful over a statistical sample, you can then run aiR for Review over the rest of your data (up to 100,000 docs at a time) and progress with your review as normal.

Altlaw Prompt-Writing Questionnaire

Answer the questions in this questionnaire to generate the material for your initial aiR for Review prompt. [CLICK HERE](#) to open the complete interactive version of this form!

1. Matter Overview:

1a. Please provide a short summary of your case.

This should be approximately 1-2 paragraphs.

Enron, an energy and commodities company filed for bankruptcy in December 2001 after its fraudulent accounting practices and financial mismanagement came to light. The company's downfall has resulted in significant legal action, including criminal charges against key executives and civil litigation involving investors, employees, and creditors.

The central issue in the case is Enron's use of deceptive accounting techniques, including special purpose entities (SPEs), to conceal debt and inflate earnings. These tactics misrepresented the company's financial health to investors and the public, artificially sustaining stock prices. Key executives, such as CEO Jeffrey Skilling, CFO Andrew Fastow, and Chairman Kenneth Lay, have been implicated in orchestrating and benefiting from the fraudulent schemes. The collapse has led to massive financial losses for stakeholders and exposed systemic issues in corporate governance and regulatory oversight.

The goal of this review is to uncover evidence of fraudulent activity, such as intentional misrepresentation of financial data, obstruction of justice, or efforts to mislead regulators and investors....

1b. Please list all your key custodians and any associated aliases.

1. Jeffrey Skilling - CEO of Enron
2. Andrew Fastow - CFO of Enron
3. Person A.....
4. Person B.....

Example Text

1c. Please list all noteworthy organisations in your case

1. Arthur Andersen LLP
2. Securities and Exchange Commission (SEC)
3. LJM Partnerships (LJM1 and LJM2)
4. Organisation 1.....

Example Text

2. Relevance:

2a. Please provide relevancy criteria in short, succinct sentences.

Documents that show the use of SPEs to offload debt, inflate profits, or manipulate financial statements are very relevant. Evidence of how these actions were designed to mislead stakeholders is critical.

Evidence of key executives engaging in fraudulent practices is relevant. Evidence where they are doing so knowingly or discussing the intent behind such actions is very relevant.

Communications or documents showing Arthur Andersen's involvement in enabling, ignoring, or failing to detect fraud are relevant. Evidence of deliberate destruction of documents by auditors is highly relevant.

Documents showing financial losses suffered by investors, employees, and creditors are relevant. Evidence that quantifies the damage or ties it directly to fraudulent activities is particularly relevant.

Documents that highlight gaps in governance, failure to meet regulatory requirements, or ethical lapses in corporate decision-making are relevant. Evidence showing systemic issues that allowed the fraud to persist is very relevant.

.....

[Fill in the complete form](#)

Prompt writing best practices

When it comes to developing your prompt it is important to remember that prompt engineering is NOT prompt writing. Writing clear prompt criteria is 90% writing clearly and will make your prompt engineering process far easier.

This means providing clear and concise instructions, using an active voice and sentence structure, avoiding confusing language (like double negatives) and avoiding the use of complex legal jargon - unless you define the terms within the prompt.

You should also run spelling and grammar checks over your prompt and ensure it is formatted in an easily readable way. Think, **“If I were giving these instructions to a human, would they understand?”**

Example:

Any and all internal written correspondence between Oct 4th 2017 and the present that relates to the manufacture of laptops, including but not limited to:

- a. Material sourcing
- b. Microchip design

Becomes...

Documents that relate to the manufacture of laptops are relevant.
Documents specifically relating to material sourcing and microchip design are very relevant.

Note: You will date filter before promoting your documents to review therefore there should be no need to specify date ranges in your prompt.

The aiR algorithm already starts with a great deal of fluency in the law but will need context on the specifics of your matter. It is also important to remember that less is more. Short, succinct sentences are far easier for the algorithm to understand than lengthy, complex paragraphs. There is also a 10,000-character limit to the total text entered into the instructions section.

Let's take a look at some more examples...

Guidance	Instead of this...	Write this...
Is this easy for a human to understand?	Industrial spying, because of the growing use of computers to store and process corporate information, is increasing rapidly.	Industrial spying is increasing rapidly because of the growing use of computers to store and process corporate information.
Reduce 'fluffy' and imprecise descriptions.	The description for this product should be fairly short, a few sentences only, and not too much more	Use a 3 to 5 sentence paragraph to describe this product.
Instead of saying what not to do, say what to do instead	Do not consider the document relevant if it only discusses project Pegasus and not project Thor	If the document discusses project Thor, it should be considered relevant. If the project discusses only project Pegasus, it should be considered not relevant.
Capitalising words can point the model towards specific instructions.	Any document which discusses jellybeans should be considered responsive.	If a document discusses jellybeans, it MUST be considered responsive.

These examples are taken from Relativity's recommendations and guidance; you can find their full list of suggestions [HERE](#).



When it comes to providing definitions to the algorithm and including certain criteria, the algorithm can be quite intuitive. For example, you need only provide the algorithm with email and nicknames for key players in your case if their email or nickname varies significantly from their given name. If a human would know who you are referring to by the nickname then so will the algorithm!

This being said, any and all uncommon jargon, codenames and abbreviations should be defined.

Ideally, you should aim to condense your Review Protocol down to about 4 pages worth of content to enter it into aiR's prompt window. So, be direct, understand that the algorithm is able to make some assumptions, and don't be afraid to capitalise to stress instructions.

So there you have it, Altlaw's complete guide to Prompt Engineering. For more information about Relativity aiR for Review check out our blogs...

[What is aiR for Review?](#)

[aiR for Review: Use Caes](#)

or speak to an Altlaw expert today to discuss how aiR can be used in your next eDiscovery project.

