

FACTSHEET

The Art of Prompting

“What are the requirements for a data processing agreement?” Questions like this, when put to AI services, often produce incomplete or overly general answers. The wording matters. Try it: “You are an experienced privacy lawyer. Analyse the requirements for a data processing agreement under the GDPR, with reference to the relevant articles.” The difference? The art of prompting.



Generative AI: a new legal assistant

Generative AI (genAI) can generate text, but it can also analyse input: spot weak points in arguments, make sense of feedback, or unpack legal issues, for example. In that role, it functions like a legal team member: drafting initial advice, analysing contracts, and developing legal arguments. It can even do this with a large text, such as a full contract or a legal article, as part of the input.

“Prompting” is the art of instructing these AI systems. Just as you would guide a new legal team member, you need to give these systems the right context and guardrails. A good prompt is like a clear instruction: specific, goal-driven, and supported by the right context. The difference between a weak and a strong prompt can be as significant as the difference between unclear and clear instructions to a colleague.

It is important to realise that these systems do not replace legal expertise. They are a tool that, if used correctly, can enrich and speed up legal work. They do not know up-to-date case law and cannot give a definitive legal judgement. Their strength lies in supporting with analysis, research, and generating first drafts. Read more in our research report: [“Can you trust AI for legal work?”](#).

Context is king

Just as a trainee lawyer performs better with clear instructions, AI systems thrive on clear context. By presenting yourself as in-house counsel seeking advice for a specific jurisdiction, you prevent the system from defaulting to general or US legal principles. A prompt such as “As a Dutch in-house legal counsel, I am looking for advice on...” immediately sets the right tone.

In general, AI systems perform better when they first see an example of the desired output format. This is called “few-shot prompting”. Instead of asking straight away, “Review these general terms and conditions”, you can first give an example: “This is what a legal analysis of general terms and conditions should look like: [example]. Now analyse these terms and conditions in the same way: [text].”

An example. Compare these two prompts for the same question:



Weak:

“What are the rules for cookies on a website?”

Strong:

“As in-house counsel for a Dutch e-commerce company, I need to update our cookie policy.

Analyse the current legal requirements under the Dutch Telecommunications Act and the GDPR, with specific attention to:

- The distinction between strictly necessary and non-necessary cookies
- Requirements for valid consent
- Information obligations in the cookie statement
- Where relevant, refer to guidance from the Dutch Data Protection Authority (AP) and recent case law.”

The difference

The second prompt defines the legal context (Dutch law), specifies the relevant legislation, and structures the desired analysis. This not only produces more targeted answers but also makes it easier to assess the quality of the advice. Where the first prompt often leads to generic marketing copy or US concepts such as “opt-out cookies”, the second prompt produces legally relevant analyses that are directly usable in Dutch practice.

One risk with the final bullet point is that the system may invent the requested guidance or case law instead of providing it. LLMs are not a search engine and do not “know” facts. Always verify this, including any specific quotations or statutory articles mentioned.

Use system prompts

Many AI services offer the option to set a “system prompt”: a kind of permanent instruction that applies to every conversation. This is, in effect, briefing your virtual legal assistant on their role and expertise. An effective legal system prompt could look like this:

“You are an experienced Dutch in-house counsel with 15 years’ experience in corporate law and privacy law. You provide practical advice, with concrete references to Dutch legislation and case law. If you are unsure about recent legislative changes, you state this explicitly. Base your answers on Dutch legal practice, not on general or US principles.”

This instruction ensures that each subsequent answer focuses on Dutch law, remains practice-oriented, and is clear about any uncertainty.

A further improvement is to enrich the system prompt with style guides or whitepapers containing the specialist knowledge you use in your work. The AI service then analyses these for style and language and applies the points raised. Think of “address the client as ‘you’” or a standard structure for advice documents.

Add structure with tags and brackets

An effective technique for structuring your prompts is to use tags or brackets. This visual separation between different parts of your prompt helps the AI system clearly distinguish input, context, and instructions. Tags work like legal headings in a contract: they clearly mark where one part ends and the next begins.

For example:

Weak:

“Here are our general terms and conditions, analyse them under Dutch law and pay particular attention to the liability clauses.”

Strong:

<context>

Dutch B2B company in the IT sector

</context>

<terms>

[Insert the text of the general terms and conditions here]

</terms>

<instruction>

Analyse the terms within the “terms” brackets, focusing on:

1. Liability clauses
2. Compliance with mandatory law
3. Comparison with industry standard

</instruction>

Brackets are particularly effective for marking examples and desired output. For instance, you can indicate:

- [EXAMPLE] for a model analysis

- [CONTEXT] for background information
- [ANALYSIS] for the desired analysis method
- [OUTPUT] for the desired format

A practical example:

[CONTEXT]

Our organisation provides SaaS services to healthcare institutions.

[EXAMPLE ANALYSIS]

A legal analysis must include the following elements:

1. Relevant legislation (GDPR, Medical Treatment Contracts Act (WGBO), Health Insurance Act (Zvw))
2. Concrete references to articles
3. Practical recommendations
4. Risk assessment

[DOCUMENT]

[Insert the data processing agreement to be analysed here]

[INSTRUCTION]

Analyse this data processing agreement in line with the format above.

This structuring has several advantages:

- The AI system receives a clear framework for the analysis
- You can easily amend or update specific parts
- The output becomes more consistent and easier to compare
- It reduces the risk of confusion between context and instruction

One key tip:

Be consistent in your use of tags or brackets within a single conversation. As with legal documents, consistency helps maintain structure and clarity.

Step-by-step analysis

Many people use LLMs as if they are Google: type a question or a few keywords and expect a complete answer. But just as you would instruct a legal team member to work in a structured way, you can guide an AI system through an analysis step by step.

Compare these two approaches to the same case:

Unstructured:

“Is an employer liable if an employee falls down the stairs while working from home?”

Step-by-step:

“Analyse the following case step by step: An employee falls down the stairs at home during working hours and suffers injury. He holds his employer liable.

Follow these steps in your analysis:

1. Identify the legal framework (relevant provisions of the Dutch Civil Code on employer liability)
2. Determine whether working from home falls within ‘in the course of employment’
3. Analyse the employer’s duty of care in a working-from-home setting
4. Review relevant case law on accidents while working from home
5. Assess any contributory negligence by the employee
6. Formulate a conclusion with practical advice
7. Include concrete references to legislation or case law at each step.”

This step-by-step approach forces the system not to skip steps. It makes the legal reasoning transparent, supports verification of sources, and prevents “tunnel vision” on a single aspect.

This method is known as “chain-of-thought prompting”. You can strengthen it further by asking the system to make its reasoning explicit:

“Before you give your final answer, first explain:

- Which areas of law may be relevant here
- Why you consider certain statutory provisions applicable or not applicable
- How different legal sources relate to each other
- Which practical considerations play a role”

An additional benefit is that the step-by-step build-up allows you to correct the system at any point if it heads in the wrong direction. As with guiding a legal team member, this prevents an initial misunderstanding from affecting the entire analysis.



Grounding in sources

Legal work is driven by sources. By explicitly asking for statutory provisions and case law, you get more specific, verifiable answers. “Analyse this situation with reference to relevant provisions of the Dutch Civil Code and recent case law” yields more usable results than a general request for analysis. A limitation for Dutch legal practice is that these systems are trained mainly on English-language and US legal sources. This creates specific risks that require a targeted approach. There is also a tendency to “hallucinate” sources, in other words, to invent them.

A solid structure, such as the framework below, can strongly improve answer quality. In addition, the iterative prompting discussed above is important for quality assurance. The more text an AI system is allowed to generate in response to a question, the greater the chance that errors creep in. A practical approach is to have the AI system first produce an outline with numbered headings and then request a detailed draft per heading (“now write section 3.2.5 and limit yourself to the blacklist of general terms and conditions”).

The framework also includes your own source material. This is a simple way to “ground” AI output, in other words, connect it to the real world. AI systems are guided more strongly by information in the input than by generally known information. For example, someone who asks “cover the black list of general terms and conditions in bullet points” gets weaker results than someone who asks “this is the black list, use it in bullet points”, followed by the text of the statutory provision (Article 6:236 of the Dutch Civil Code).

This systematic approach helps you mitigate the inherent limitations of AI systems and arrive at more reliable legal analyses. It remains essential that a legal professional who knows Dutch legal practice validates the output.

<legal_framework>

Applicable law: Dutch law

Relevant legislation:

- Dutch Civil Code (in particular Articles 6:231-247 DCC)
- Telecommunications Act (in particular Chapter 11)

Dutch case law only, from 2020 onwards

</legal_framework>

<verification_protocol>

For each legal conclusion, state:

- Level of certainty (certain/probable/possible)
- Basis for the conclusion (legislation/case law/doctrine)
- Any knowledge gaps or uncertainties

</verification_protocol>

<source_material>

[Copy relevant statutory provisions]

[Key passages from case law]

[Quotations from authoritative literature]

</source_material>

<instruction>

Base the analysis exclusively on the sources above.

Mark where you deviate from these sources.

</instruction>

<quality_criteria>

- Only verifiable source references
- Explicit distinction between facts and interpretation
- Marking of knowledge gaps
- Indication of how current the sources are

</quality_criteria>

Use your own knowledge as a foundation

A frequently overlooked way to get an LLM to produce better answers quickly is to feed it knowledge from your own documents. Your firm or department has piles of whitepapers, analyses, advice notes, or perhaps handbooks that colleagues rely on. Why not let the LLM do the same?

Compare these two scenarios:

Without your own knowledge:

“Assess whether this liability clause is valid...”

With your own knowledge:

“Here is our current liability clause: [text]. This is what my colleague Michael recently blogged about limitation of liability clauses in SMEs: [text]. Here is relevant case law from the Amsterdam Court of Appeal in our sector: [quotation]. And this is what the industry association states in its model terms: [text].

Analyse to what extent our clause:

1. Aligns with current case law
2. Deviates from the industry standard
3. Requires amendment”

The difference? The second prompt gives the system concrete and relevant context to work with, which leads to more targeted and practically usable answers. The use of tags is crucial here to keep the different source texts separate.

The major AI providers now offer ways to permanently feed the system with your legal knowledge. This effectively creates a personalised legal assistant that knows your organisation inside out. The key is to give the system enough context to be valuable, without overwhelming it with irrelevant information. Start small, for example with frequently used templates and standard terms, and then build from there based on practical experience.

Microsoft's Copilot

Offers the option to select a SharePoint folder with relevant documents (contracts, memos, procedures) from which Copilot can draw when answering questions. The strength of this is:

- Automatic integration with SharePoint and Teams
- Real-time updates when documents change
- Preservation of existing access rights
- Ability to define specific document sets per question type

OpenAI

Calls this feature “Custom GPTs”. This allows you to create a personalised version of ChatGPT. For a legal department, this means:

- Upload your own legal documentation
- Define specific instructions and working methods
- Build in templates and example prompts
- Ability to share directly with the team

Many Custom GPTs are public. They can also be made accessible by invitation, and within an Enterprise subscription they can be made accessible only to other users within the organisation.

Anthropic's Claude

Claude offers a different approach with “Projects”. Instead of permanent knowledge storage, you can upload relevant documents per conversation.

Advantages:

- More control over what information is shared and when
- Flexibility to use different sources for different questions
- No risk of knowledge spill-over between projects
- Better privacy safeguards

Within the “Claude for Work” subscription, access to projects can be restricted to colleagues only.



Closing: the art of effective prompting

Let us briefly recap. The core principles we discussed:

1. Structure and context

- Use tags and brackets to separate different components
- Provide relevant legal and business context
- Define the applicable legal framework
- Make expectations about source references explicit

2. Step-by-step analysis

- Break complex issues into sub-steps
- Have the system make its reasoning explicit
- Validate interim conclusions
- Adjust where needed

3. Anchoring in sources

- Provide the system with concrete statutory texts
- Cite relevant case law
- Integrate your own legal documentation
- Verify all source references

Additional practical tips:

Iterative refinement

- Start with a concise question to explore the direction
- Refine based on the first answer
- Ask follow-up questions on specific parts
- Use the output as a basis for further questions

Template development

- Develop standard formats for common analyses
- Build a library of effective prompts
- Document successful prompting strategies
- Share best practices within your organisation

Quality assurance

- Have the system explicitly state where it is unsure
- Ask for substantiation for critical conclusions
- Spot-check source references
- Document any inaccuracies found

Knowledge integration

- Gradually add more of your own knowledge
- Test new prompts first on known cases
- Evaluate the added value of different source types
- Monitor the impact on answer quality

By applying these principles and tips consistently, you develop an increasingly effective dialogue with AI systems. The result: sharper analyses, more reliable output, and more efficient legal workflows. Keep experimenting, document your experiences, and share insights with colleagues. That is how your expertise grows in this new way of doing legal work.

Are you concerned about the GDPR, confidentiality, or (professional) liability when using AI professionally?

Our research report [“Can you trust AI for legal work?”](#) also covers this in more detail.





**Would you like to
receive additional
information?**

Want to know more?
Get started with AI?
Contact us at:

e-mail: contact@ictrecht.nl

telephone: 020 663 19 41.