

Doctrine Code of Conduct on Generative AI

Why did we adopt this Code of conduct ?

At Doctrine, we have been using artificial intelligence since our beginnings to provide our legal intelligence platform to law professionals. The integration of generative artificial intelligence raises new and legitimate questions that we have sought to address transparently for our users. This code of conduct on generative AI is a unilateral commitment to adhere to the following principles in our innovations and developments integrating generative artificial intelligence.

Why use Generative AI ?

Generative AI enables significant advancements in natural language processing, which is at the core of Doctrine's technology. Through large language models (LLMs) such as GPT-3.5, GPT-4, or Mistral models, we can now accomplish certain tasks more quickly and/or efficiently. This will allow us to build features that will save even more time for our users.

How does Generative AI work ? What are the advantages compared to other solutions ?

Generative artificial intelligence utilizes large language models to generate responses based on user queries. At Doctrine, we place ourselves at the intersection between (1) your legal inquiries and needs, (2) generative AI models, and (3) legal information from our database. We differentiate two main categories of use-cases: those without user input ("prompt") and those requiring it.

For **prompt-less use-cases**, we identify, qualify, and analyze a need, and conduct business analyses to transmit this need to our generative AI providers anonymously and securely. For instance, summarizing decisions or documents can be done with a single click, without transmitting the identity of the user requesting the summary, and with pre-processing by Doctrine's experts.



For **prompt-based use-cases**, we first work on understanding your question to determine where to search. We retrieve relevant legal information from our database to source our response. Then, we transfer your processed question and this retrieved information to our generative AI providers to deliver an answer enhanced by our mastery of legal information contextualization (what we call the Doctrine Legal Graph). In essence, we commit to bringing added value to interactions with generative AI providers and not merely acting as a middleman for your questions and needs.

What are your guarantees in terms of data security ?

Security principle

Law professionals handle highly sensitive and confidential data, and we understand that we must earn your trust when you use our services.

That's why Doctrine heavily invests in the security of all data, whether personal or not, sensitive or not, that passes through our platform. We are aligned with the [ISO 27001](#) standard and currently undergoing certification.

We apply our security policy to all operations we perform, which includes strict internal organizational rigor, encryption of all data in transit and at rest, and hosting with providers recognized for their security (our AWS provider is certified ISO/IEC 27001:2022, 27017:2015, 27018:2019, 27701:2019, 22301:2019, 20000-1:2018, 9001:2015, and CSA STAR CCM v4.0).

Furthermore, we follow best practices in terms of development and deployment to ensure that we do not introduce vulnerabilities during our iterations, including automated testing, human review of each new feature, and separation of development and production environments.

To ensure that our security complies with current state-of-the-art, we also invest in the services of two major providers : Datadog, which allows us to combat certain forms of automated attacks and monitor platform usage to detect anomalies, and Yogosha, which allows us to have our security tested by ethical hackers whom we reward if they discover vulnerabilities (penetration tests).

If you would like to learn more about our security policy, please do not hesitate to contact us via email at dpo@doctrine.fr.

What are your guarantees regarding the confidentiality of personal data ?

Confidentiality principle

The subject of personal data is particularly significant at Doctrine. Our first innovation, in 2016, was an algorithm for anonymizing judicial decisions, [to reconcile privacy and public data](#). Our commitment to GDPR compliance has allowed us to successfully pass an inspection by the CNIL. This commitment does not waver with the production of innovations using generative artificial intelligence :

- We do not transfer identifying data to our generative artificial intelligence providers. This means that we only send legal queries, without associated identifying metadata, containing only the result of our research on public data and what we have understood from your query, reformulated. This commitment is essential as it allows for a complete separation of the “response generation” part from the previous understanding of the question and searching for relevant sources, which takes place entirely within our architecture.
- We continue to store your data in Europe only. For the use in production of generative AI features with prompts, we have obtained access to the GPT service from Microsoft Azure Europe. This means that we use GPT 3.5 and GPT 4 models but within an Azure environment, located in Europe, without transferring data to the United States. It is worth noting that [the Conseil d'État considered](#) that these guarantees allowed the Health Data Hub to host its data in a similar manner. We will only store your data outside of Europe in cases of absolute necessity, informing you accordingly, and providing all necessary legal and technical guarantees to prevent any risk of access to transferred information.
- We have our providers sign the standard contractual clauses of the European Commission, even if they are located in countries that do not require it, in order to contractually support our technical commitments.

More generally, our general terms and conditions have contained since a long time the following confidentiality commitment, which applies to our generative AI services : Doctrine irrevocably undertakes, except in cases of legal obligation or injunction by judicial authority, to respect the strictly confidential nature of the data generated by your use of its services. Doctrine prohibits itself and its employees from using, allowing the use of, or disclosing to anyone, directly or indirectly, in any form whatsoever, the content of this data. Doctrine acknowledges that a breach of this commitment may result in its liability.

This means, in particular, that unless expressly agreed by you, we do not reuse your data for the purpose of training our artificial intelligence systems.

How can I tell that you are using Generative AI?

Transparency principle

Transparency is at the core of digital regulation and platforms, which is why we have chosen to integrate it as a fundamental principle of our innovations using generative AI. This commitment is

demonstrated by this code of conduct, as well as by another simple and effective guarantee : we commit to informing you whenever we display content generated by artificial intelligence.

More generally, we also commit to being transparent about the providers we use in our personal data policy, allowing you to know at any time what we are doing and with whom. Finally, we commit to being transparent about the capabilities (and thus the limitations) of our features, so that you can use them with full confidence, having mastered all essential information.

Am I required to use your Generative AI features ?

Free choice principle

The answer is simple : no. And for two reasons.

First, because naturally, you can always choose not to use the features available on the site.

But to take into account the various sensitivities regarding these latest innovations, we have developed an opt-out system for prompt-based generative AI features, which blocks them for you and your organization. If necessary, do not hesitate to contact your account manager to enforce it. This opt-out does not entitle you to any refund or reduction.

What responsibility do you take for your Generative AI features ?

Shared responsibility model

On this topic, as on all topics related to the use of legal intelligence platforms by legal professionals, our approach is one of **shared** responsibility

Indeed, as a law professional, especially as a lawyer, judge, notary, or accountant, you remain subject to your own code of ethics, and the use of a professional platform such as Doctrine cannot exempt you from it. Any use, not only of our generative AI features but in fact of any legal information provider, is thus done under your responsibility. Therefore, you must verify any important information.

However, we do not believe in a model aimed at completely absolving legal intelligence platforms from responsibility. We cannot guarantee the accuracy of the content produced through our generative AI features because, even though they are based on our sources, errors are possible (as with any content produced by a human). However, we commit to adhering to this code of conduct and all its principles. We also strive to minimize erroneous generations (“hallucinations”) by

consistently providing the necessary context to our generative AI providers, with the understanding that hallucinations cannot be reduced to zero, contrary to what some industry players claim.

What is your stance on copyright and trade secrets ?

Respect for pre-existing rights principle

Generative AI raises interesting questions regarding copyright, both regarding model training and ownership of generated content. Most of these questions are currently being studied by legislators and regulators, nevertheless we wish to make the following commitments :

- We respect copyright in the implementation of our generative AI features. This notably means that we do not transfer the full content of copyrighted articles to our generative AI providers, unlike some industry players, without the consent of the rights holders or exceptions to copyright.
- We respect your copyright and trade secrets by preserving the confidentiality of your data and not reusing it for the purpose of improving our services, except with your explicit consent.
- We do not claim any rights over the generated data, which you are free to reuse, and we commit to only work with generative AI providers who align with this position.

What are your commitments in terms of ethics ?

Ethics by Design principle

Doctrine has a legal and compliance department led by a qualified lawyer, and ensures to integrate compliance and ethical considerations from the outset of all its products. This is why we are already signatories of the [Charter on transparency and ethics in the use of judicial data](#) by the National Bar Council and are [certified "PrivacyTech"](#) by the AFNOR.

In this context, we commit to :

- Always involve a member of the legal team in the design and development of our features using generative artificial intelligence. In practice, our legal team regularly contributes to these features by providing legal expertise.
- Work to limit biases in our training and testing datasets.
- Regularly audit our innovations in production to ensure that their behavior remains compliant with requirements.

Of course, these commitments will be strengthened by our compliance with the European regulation on artificial intelligence, following its promulgation by the European Commission.

Do you collaborate with the legal ecosystem on these subjects ?

Virtuous innovation principle

The subject of generative artificial intelligence in the legal sector cannot be fully grasped by a single actor. Therefore, we commit to participating, as much as we can, in initiatives brought to our attention to address these issues and achieve a general consideration of the challenges raised and the solutions to be provided. We will participate in such initiatives, particularly within the Legaltech group of France Digitale, which we co-lead.