

# VLAM legal derisking analysis

## Executive summary

13 May 2026

### 1. General

This executive summary relates to a derisking analysis that was commissioned in response to the growing strategic importance of generative AI and the parallel increase in legal and regulatory exposure associated with its use within the VLAM research project. The analysis was drafted by Jozefien Vanherpe en Louise de Béthune (KU Leuven CiTiP), in Dutch.

This derisking analysis is grounded in a European context. Operating within the EU is a strength. It provides legal certainty, credibility, and trust, which are essential for long-term value creation in AI-driven ecosystems. At the same time, the EU framework comes with clear legal and regulatory requirements. These rules are not optional, and ignoring them would expose the VLAM project and its partners to material legal, financial and reputational risks. The purpose of this analysis is to make those requirements explicit, not to discourage innovation, but to enable it in a way that is sustainable and defensible.

The analysis is deliberately structured around a worst-case scenario logic. By mapping the full spectrum of potential legal risks across the AI lifecycle, from input to output, it provides a first, robust step toward informed decision-making and risk prioritisation. This approach is intentional. It is far easier to innovate responsibly once the outer boundaries of risk are well understood. The document should not be read as an argument to stop data sharing, halt AI development or freeze experimentation. It is not intended as a final verdict on technical or commercial feasibility, but as a foundation for further, more targeted analysis, and for governance choices that balance legal certainty with technological progress.

This analysis should be read as an enabler for new EU-based AI projects. Europe's regulatory framework sets the rules of the game, but it does not prohibit innovation. On the contrary, future economic growth and societal welfare depend on accelerating responsible innovation in AI. The real danger lies in overreacting to the identification of legal risks by freezing development or avoiding data use altogether. By understanding the worst-case scenarios upfront, VLAM and its partners are better positioned to innovate faster, with confidence, and on solid legal foundations.

## 2. Structure

The analysis starts by establishing a common legal baseline (**section 1**). It outlines the essential principles of intellectual property (IP) law that are most relevant for generative AI, focusing on copyright and the protection of trade secrets. These areas are most directly affected by the use of large datasets, the training of models and the generation of VLAM AI technology (in software form) and, at a later stage, new content. Other IP regimes, such as trademark and patent law, are not ignored, but deliberately left out of the core analysis to keep the focus practical and relevant. They can be addressed later if specific commercial or technical developments make them relevant. In parallel, the analysis recalls the fundamentals of European privacy and data protection law, with the GDPR as the central reference framework, focusing on data protection by design and by default, as well as measures to comply with data subject rights (right to information, right to erasure, right to rectify personal data). It also introduces the core logic of the AI Act, which will increasingly shape how AI systems must be designed, documented and governed in the EU. Finally, it looks at platform liability rules that may apply to (gen)AI interfaces.

Against this background, the analysis follows the lifecycle of AI systems used or developed within the VLAM project. In the input phase, where AI models are trained (**section 2**), the main legal sensitivities relate to the origin and nature of the data. From an IP perspective, the use of copyright-protected works or protected databases without a clear legal basis can create exposure. The analysis distinguishes different types of content included in the training data and links them to the relevant IP right(s). In this context, focus not only lies on the scope of protection of such IP right(s) (including both exploitation rights and moral rights), but also the validity requirements and issues relating to the titularship of such IP right(s). From a data protection perspective, the presence of personal data triggers responsibility and obligations under the GDPR. Personal data, including sensitive data, must be clearly identified, to mitigate the risk of data breaches and complaints. The AI Act adds governance and risk-management expectations. These issues are not raised to suggest that training AI models in the EU is impossible. Rather, they show that training must be organised deliberately, with conscious choices about data sources, documentation and safeguards.

The analysis then considers situations where the project relies on external or third-party AI models (**section 3**). Here, the critical question is not only how well a model performs technically, but whether it can be used freely and safely from an IP and contractual point of view. License terms, usage restrictions and downstream obligations can significantly limit freedom to operate (FTO) if not assessed in advance. This reinforces the broader point that legal and technical architecture are inseparable when scaling AI solutions.

In the output phase (**section 4**), attention shifts to what AI systems produce. Different types of output give rise to different risks. Some outputs may conflict with existing rights held by the media partners themselves, while others may affect the rights of third parties. At the same time, the project must consider how newly generated outputs are protected, owned and managed as foreground IP. These questions directly affect commercial exploitation, strategic partnerships and long-term value creation. Again, the analysis presents the highest-risk scenarios first, to make informed prioritisation possible.

The final synthesis (**section 5**) brings together these findings in a structured risk overview and a plan of approach. The key message is that not all risks require the same response. Some can be mitigated through governance, contracts and internal processes. Others require strategic decisions about scope and ambition. What matters most is that these choices are made consciously and aligned with the project's broader goals.

### **3. Link with other research projects**

The legal risks identified in the VLAM study are not conceived as static or permanent obstacles. On the contrary, several ongoing and future research projects explicitly build on this groundwork with the objective of operationalising solutions, reducing uncertainty, and embedding compliance directly into technical and organisational practices. Together, these projects ensure that the issues identified in the VLAM analysis will evolve from abstract legal risks into manageable, technically supported design constraints.

A first and central continuation of this work can be found in the ARCHAI project. ARCHAI directly leverages the legal risk matrix and action plan developed in the VLAM study as a foundation for further research. In particular, the findings of Chapter V of the VLAM study identify which legally relevant information points within media data are crucial for AI training, such as information linked to copyright status, licensing conditions, personal data, and AI Act-related obligations. This preliminary legal mapping allows ARCHAI to move beyond problem identification and to focus on legal modelling. The project aims to encode regulatory and contractual constraints into machine-actionable metadata, thereby translating legal requirements into technical artefacts that can be enforced automatically throughout AI pipelines.

In parallel, the ART-AI project provides complementary input at a more conceptual and systemic level. ART-AI focuses on analysing existing policy languages and vocabularies used to track legal rights in data. Its work will offer an up-to-date overview of current solutions, while also identifying gaps, overlaps and potential synergies across different approaches. ART-AI's final report on best practices for deploying automated compliance technologies for EU digital legislation will serve as a direct input for ARCHAI.

By aligning ARCHAI's legal modelling efforts with ART-AI's comparative analysis, unnecessary duplication of preparatory work is avoided and design choices can be grounded in a broader European state of the art.

Further legal and technical insights will emerge from the cSBO project MUMIA, which focuses on multimodal, agentic AI systems in a media context. The legal research within MUMIA specifically addresses how media organisations can develop AI systems capable of generating high-quality, long-form video content based on Flemish audiovisual archives, while remaining legally compliant. This work directly engages with many of the issues identified in the VLAM study's output phase, including background IP, reuse of archival content, and the lawful generation of derivative works. As such, MUMIA can be seen as a practical testbed for the legal principles and risk assessments previously identified.

Finally, the ICON project DAGA will address one of the most sensitive and societally relevant challenges associated with generative AI: attribution and remuneration. With a dedicated legal pillar, DAGA will develop technology that can determine which music, images or textual sources contribute to AI-generated outputs. Its objective is to enable accurate acknowledgment and fair remuneration of artists and rightsholders. This work directly responds to concerns raised in the VLAM study regarding output-related infringement risks and the governance of foreground IP, by offering concrete technical mechanisms to support legally and ethically acceptable reuse.

Taken together, these projects demonstrate a coherent research trajectory. The VLAM study identifies and structures legal risks as a necessary first step. ARCHAI, ART-AI, MUMIA and DAGA then progressively translate those risks into models, tools and governance mechanisms. Rather than remaining problematic, the issues identified in the VLAM analysis become drivers for innovation in legal-by-design AI systems, reinforcing the broader ambition to innovate faster while staying firmly anchored in the European legal framework.

#### 4. Risk matrix

Rank	Name	Description	Impact (1-5)	Likelihood (%)	Mitigation	Effort (1-3)
1.	<b>output: liability</b>	liability for infringing output, both © (incl. ° summaries) and TM law (e.g. ads used as input) + GDPR (data breaches)	4	85	guardrails (output constraints), memorisation tests, snippet limitations, no verbatim reproduction, content controls for TM (linked to TM databases?), T&C, content moderation, notice and take down policy	3
2.	<b>input: training 1</b>	GDPR infringement through training process (incl. special categories) + action by DPA(s)	5	60	privacy by design: metadata annotation (DPV), lawful basis, purpose-limitation and data minimisation, anonymisation, DPIA	3
3.	<b>input: training 2</b>	qualification as (1) 'communication to public' / (2) 'reproduction' under © law + express opt-out by artist(s)/ CMO(s) (speech-to-text process, annotation process, training process)	5	45	chain of title (link with annotation process), metadata annotation, explicit warranties by partners, model clauses/addenda for partners' contracts	3
4.	<b>input/ output: data subject rights</b>	request(s) for access, erasure and objections to processing	4	50	privacy policy, complaints mechanism (including a takedown policy: access limitation for the future), considering public interest	2
5.	<b>output: ownership 1</b>	dispute(s) re ownership + exploitation of model(s) (who can do what?) + AI	3	15	co-ownership partners = shared responsibility, reciprocal license	1

		Act responsibility (who is 'provider?')				
6.	<b>input: existing models</b>	limitations to commercial use of software built upon	3	<i>depends</i>	choose permissive models without commercial use restriction, e.g. Apache 2.0	1
7.	<b>output: systemic risk</b>	if FLOPs are greater than $10^{25}$ the model qualifies as AI model with systemic risk with extra obligations (unless argued)	2	80	notification and argumentation / potential risk management	1
8.	<b>input: transparency</b>	AI Act: IP policy, technical documentation ...	2	70	transparency and information incl. machine-readable	1
9.	<b>general</b>	evolution of law: difficult to predict + novelty of legal analyses	2	60	continuous assessment over time	1
10.	<b>input: training 3</b>	content where initial rights holder unclear	1	80	best efforts use of orphan works regime	1
11.	<b>output: ownership 2</b>	dispute(s) re effect of lack of AI authorship/ownership on output ownership/exploitation	2	40	T&C, (sub)licenses	2
12.	<b>output: TS</b>	output reveals trade secrets (/ info that partners would seek patent protection for	2	5	guardrails (cf. 1)	3