Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 134 (2025) 999–1004

www.elsevier.com/locate/procedia

Proceedings of the 58th CIRP Conference on Manufacturing Systems 2025

Legal Implications of Vision-Language Foundation Models (VLFM) in Industrial Applications in Europe: An Inquiry into Data Protection, Copyright, and AI Regulation

Keno Moenck[*,a], Niklas Wais[b], Martin Gomse[a], Thorsten Schüppstuhl[a], Boris Paal[b]

[a] *Hamburg University of Technology, Institute of Aircraft Production Technology, Denickestraße 17, 21073 Hamburg, Germany*
[b] *Technical University of Munich, School of Social Sciences and Technology, Law and Regulation of the Digital Transformation, Arcisstr. 21, 80333 München, Germany*

* Corresponding author. Tel.: +49-40-42878-3341 ; fax: +49-40-42731-4551. *E-mail address:* keno.moenck@tuhh.de

Abstract

The rise of Foundation Models (FM) starts to revolutionize different aspects of machine vision in industrial applications. Increased generalizability, robustness, reliability, and the capacity for few-shot learning make vision systems adaptable and flexible to an unprecedented degree. Yet, the development and deployment of fine-tuned models in the industrial domain face unique legal challenges in Europe that arise from stringent regulations, including Copyright Law, the General Data Protection Regulation (GDPR), and the recently enacted AI Act. In this work, we investigate the legal implications of pre-training (1), fine-tuning (2), and utilizing VFMs (3) in industrial contexts in Europe, with a specific focus on manufacturing. With this contribution, we aim to raise awareness on how various regulations will impact the use of VLFMs in industrial settings.

Keywords: Legal Implications; Foundation Models; Manufacturing; Europe; AI Act

1. Introduction

Foundation Models (FM) fundamentally shifted Machine Learning (ML) from single-task and single-modality towards multi-task and multi-modality learners. Using model architectures that scale the training process efficiently as well as self-supervised learning methods, which allowed for non-human annotated training data, the capabilities of neural networks were significantly extended [1]. Following the success of Large Language Models (LLM), Vision and Vision-Language Foundation Models (VFM/VLFM) are trained solely on images or image-text pairs and enable a variety of different tasks related to visual perception. As an example in machine vision, traditional methods depended on explicit handmade rules, while learning-based methods, at first, made the description of how a task is solved redundant; it is implicitly induced rather than explicitly stated (s. Fig. 1). Afterwards, deep learning architectures also abstracted feature engineering by directly learning higher-level representations from samples. Following the ML paradigm that more trainable weights and broader training data increase a model's generalizability and performance, FMs are pre-trained at scale on massive amounts of training samples. Subsequently, they only need to be adapted with a few training samples for a specific domain and task or even provide sufficient zero-shot performance (i.e., do not need extra samples).

Machine vision-enabled visual perception in industrial applications is often used to increase automation and, therefore, productivity [2]. Specific tasks include assessing a product's or resource's position and orientation, shapes, dimensions, properties, and other qualitative or quantitative characteristics. In all these tasks, a vision system's adaptability and flexibility are challenged when facing decreasing lot sizes, increased uncertainty (e.g., human involvement), or increased process, product, or resource variants [3]. Most FMs meet these challenges well due to their increased generalizability and data-efficient adaption strategies, e.g., fine-tuning[1]. This results in improved robustness, reliability, and decreased engineering expense of vision systems. The core of self-supervised pre-training at scale is the amount, variety, and broadness of the training data as well

[1] We define fine-tuning as the process of (1) (partly) further training a pre-trained model's parameters or (2) adding additional trainable parameters while freezing the pre-trained model, e.g., adapter-tuning.

2212-8271 © 2025 The Authors. Published by Elsevier B.V.

10.1016/j.procir.2025.02.223

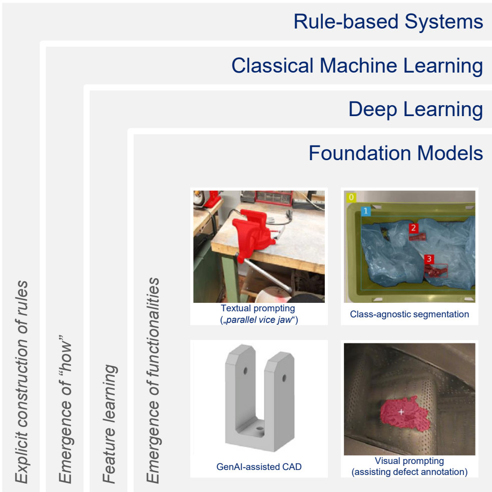

Fig. 1. Levels of intelligent systems and the enabling paradigm.

as the increased necessary computational resources. While the first is often addressed via web-crawled data, the second demands increased investment costs, which is only economically reasonable for models that can be later utilized in a considerable amount of downstream tasks across companies and even domains. Yet, pre-training on vast amounts of data and sharing pre-trained models with other parties raises various legal questions. In this work, we cumulatively elaborate on the legal implications of VFM and VLFM in industrial applications:

- We explore the permissibility of compiling large multimodal datasets for pre-training "from scratch" within Europe, examining key legal issues surrounding data acquisition from web crawling and other sources (▹ Sec. 3.1).
- We investigate the legal considerations necessary when fine-tuning existing models for specific industrial tasks (▹ Sec. 3.2).
- We assess the use of pre-trained models in industrial applications, analyzing how the AI Act influences the legal landscape for deploying these models (▹ Sec. 3.3).

This work focuses explicitly on FMs rather than deep learning techniques. The key characteristics that differentiate them from previous advances are the task-open large-scale pre-training and the necessity of fine-tuning – adapting it to the particular task and domain. In Section 3, even if we only explicitly address VLFMs, the implications usually also account for VFMs. Also, we limit our scope to the text-image modalities and do not consider other modalities, e.g., speech or music. From our perspective, they are of the highest relevance for industrial applications, even though the findings are at least partially transferable.

2. Fundamentals

2.1. Vision-Language Foundation Models (VLFM)

Machine vision tasks have significantly evolved over the past few years from simple classification, e.g., in the form of number or character recognition, to tasks that span over semantic, instance, and panoptic segmentation, object recognition, pose estimation, or sentiment analysis. With self-supervised training came the inclusion of broader language, which led to more capable caption generation and also multi-modal tasks like image-to-text and image-text-to-text models for Visual Question Answering (VQA) or Document Question Answering (DQA). Meanwhile, all major AI companies are offering pre-trained VLFMs, among them ChatGPT-V [4] and Llama-3.2-Vision [5]. While the weights of most models are publicly accessible, the pre-training datasets are often kept secret. With so many VLFMs available on various platforms, selecting a suitable model for a specific application becomes challenging - not only because details about pre-training datasets are typically sparse but also because benchmarks often cover only a limited range of real-world applications.

2.2. VLFMs in Industrial Applications

Industrial activities aim to output, maintain, and dispose of goods as well as tangible and intangible assets, which, in the case of tangible ones, can include processes from initial planning, design, production control, manufacturing, assembly, distribution, maintenance, retrofit, disassembly, and disposal.

Starting with the rise of FMs for Natural Language Processing (NLP), multiple works discuss the capabilities of off-the-shelf LLMs for industrial applications. Especially in robotics [6] and design engineering [7], promising areas arise by including text-to-text models. Here, adapting to the specific domain and task can be done inexpensively through in-context learning, compared to fine-tuning. However, current use cases often focus on assisting humans in engineering tasks rather than employing models as autonomous agents.

As of now, few works describe specific applications or survey the use of VLFMs in industrial contexts. This is probably due to the more recent developments in the vision-text area. Nevertheless, different industrial activities can involve vision-related tasks and can fully or semi-automate processes or be involved in more human-centric applications, e.g., design engineering. [3, 8] present and evaluate particular use cases of deploying a VFM while [9, 10] explore model types that include language modality. The papers do not account for legal implications, which can alter real-world use cases in the EU.

2.3. European Union (EU): Types of Legislation

All actions undertaken by the EU are based on treaties among its member states. These treaties outline the Union's goals, establish the rules governing EU institutions, define the decision-making processes, and clarify the relationship between the EU and its members. In the context of EU law, treaties represent the foundation and are referred to as *primary law*. The legal framework derived from the principles and objectives set forth in these treaties is known as *secondary law*. Legislation within the field of secondary law can be classified into several distinct categories that affect different levels of a member state's legal system. The most fundamental types are *regulations* and *directives*, some of which are highly relevent for the development and application of VLFMs in the industry (s. Fig. 2).

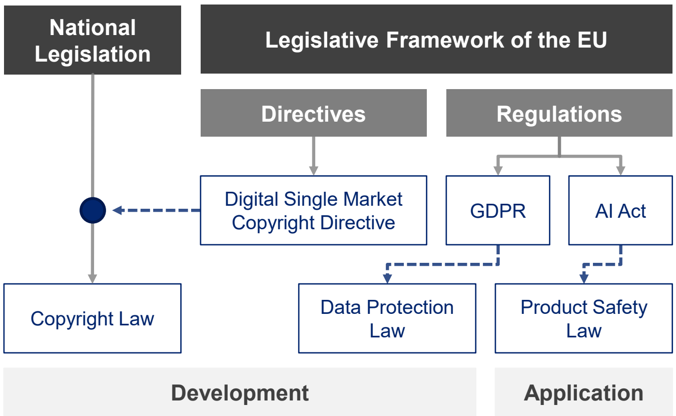

Fig. 2. Relevant directives, regulations, and laws in the scope of this work.

Regulations are legislative acts that are applicable directly in all EU member states, eliminating the need for national implementations. They are typically employed in legal areas that require consistent handling across the EU. Particularly relevant in the context of VLFMs are the recently enacted AI Act (Regulation 2024/1689) [11] and the GDPR (Regulation 2016/679) [12].

Directives have to be implemented into national laws by the member states. The vertical procedure allows for national discretion while harmonizing legal standards across the EU. Typically, directives only address particular issues in one area of the law. For the training of VLFMs, the Directive on Copyright in the Digital Single Market (Directive 2019/790) [13] is of interest.

Application. When national courts apply regulations or national laws affected by directives, they *can* ask the European Court of Justice (ECJ) for a ruling on matters of interpretation or validity. To ensure uniform application, courts against whose decisions there is no judicial remedy generally *must* do so.

2.4. AI Act

The AI Act imposes obligations on AI providers and organizations using AI systems professionally. It follows a risk-based approach, where the obligations depend on the abstract risks emanating from the product (s. Fig. 3). *Unacceptable-risk AI*, like social scoring or manipulative systems, is banned (Art. 5). Most responsibilities fall on providers of *high-risk AI* systems (Art. 6 et seq.) if they are based in the EU or their system is used within the EU. Professionals deploying high-risk AI have fewer obligations but must still comply with certain regulations if they use respective systems within the EU. *Minimal-risk AI* (Art. 4, 95), such as spam filters, remains mostly unregulated. Furthermore, there are product-specific regulations. If an AI system is intended to interact with humans, for instance, providers must ensure that users are aware of the AI (Art. 50 (1)). For systems that generate content, they have to ensure that output is marked as generated in a machine-readable format (Art. 50 (2)).

Providers of *general-purpose AI* (GPAI) must provide technical documentation, implement a policy to respect copyright rules, and provide a summary about the training data (Art. 51 et seq.). GPAI models with an open license only need to meet copyright and data disclosure requirements unless they pose a systemic risk. A systemic risks is presumed when the cumulative amount of compute used for training is greater than 10^{25} FLOPs. Any GPAI considered a systemic risk, open or closed, must undergo testing, report incidents, and ensure cybersecurity measures. Non-compliance with the obligations of the AI Act can result in fines up to 7% of global turnover or €35 million (Art. 99).

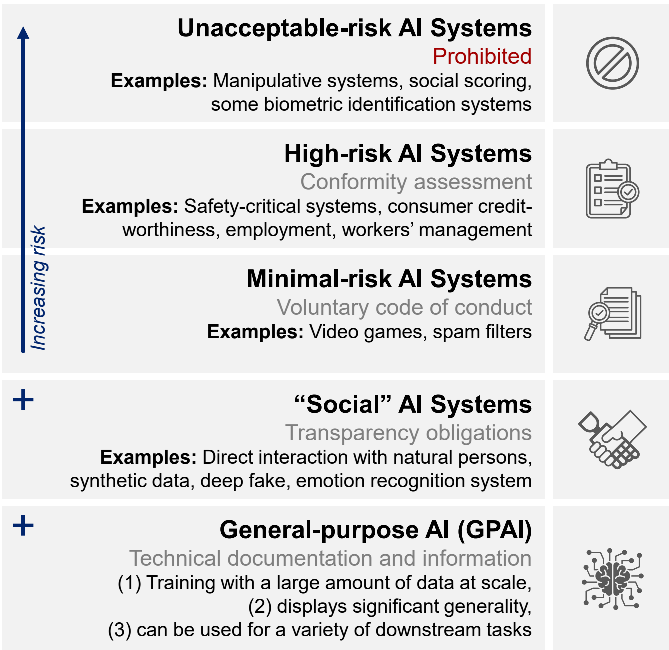

Fig. 3. Risk-based classification of AI systems and certain AI systems, which we characterize as "social" (Art. 50) and General-purpose AI (GPAI).

2.5. General Data Protection Regulation (GDPR)

The GDPR governs the processing of personal data. It applies to companies with a presence in the EU and outside the EU that operate in the European market and process the personal data of EU citizens (Art. 2, 3). Personal data refers to "any information relating to an identified or identifiable natural person" (Art. 4 (1)), including names, e-mail addresses, photographs, and online behavior. Processing means every operation that is performed on personal data (Art. 4 (2)), most notably in its collection, storage, and sharing. Personal data must only be processed on a legal basis, which can be explicit, informed consent, contractual necessity (the processing required to fulfill a contract with the individual), legal obligations (there is a legal requirement to process the data), or legitimate interests (where the processing has to be appropriate, necessary, and balanced against the rights of the individual), among others (Art. 6). Non-compliance can be sanctioned with fines up to 4% of global turnover or €20 million (Art. 83).

2.6. Directive on Copyright in the Digital Single Market

In contrast to AI safety and data protection law, which are governed directly by EU regulations, copyright law is a matter of the member states. However, the Directive on Copyright in the Digital Single Market has harmonized certain aspects. Most importantly, it mandates member states to implement copyright exceptions for text and data mining, defined as ”any automated analytical technique aimed at analyzing text and data in digital form in order to generate information which includes but is not limited to patterns, trends, and correlations” (Art. 2 (2)). Reproductions and extractions of lawfully accessible works have to be made legal ”as long as the right to do so has not been expressly reserved by their rightsholders in an appropriate manner, such as machine-readable means” (Art. 4 (3)). In the case of research organizations, however, such a reservation shall not preclude the copyright exception (cf. Art. 3).

3. Legal Implications

Concerning the legal requirements for VLFMs in the EU, a distinction can be made between developing and deploying models. During development, which typically involves the compilation of large multi-modal datasets for model training, data protection law (GDPR) and copyright law are of interest (s. Sec. 3.1). According to the current state of the debate, they also come into play when pre-trained models are used (s. Sec. 3.2). While copyright law is directly relevant during web crawling, crawling internal databases, which often include customers' data, can be directly prohibited through initial contracts on the job that produced the respective data. Regarding the practical use of the trained models in industrial applications, product safety law (AI Act) has to be considered (s. Sec. 3.3).

3.1. Compilation of Multi-modal Datasets

Datasets have to be compiled both for training VLFMs entirely ”from scratch” and fine-tuning them. If data is collected internally, data protection law becomes relevant (s. Sec. 3.1.1). In industrial applications, contractual obligations may also limit the usability of internal data; this, however, is a question of individual contractual relationships. If publicly available data is crawled from the internet, which is typically done for the task-independent pre-training of VLFMs, copyright law can impose additional limitations (s. Sec. 3.1.2). The data-related obligations imposed on providers of GPAI models by the AI Act have already been pointed out (s. Sec. 2.4). GPAI refers to a model ”that displays significant generality and is capable of competently performing a wide range of distinct tasks [...] and that can be integrated into a variety of downstream systems”. Consequently, VLFMs will generally constitute GPAI. It should, however, be kept in mind that the definition in Art. 3 (63) excludes ”AI models that are used for research, development or prototyping activities before they are placed on the market.”

3.1.1. GDPR

The GDPR comes into play once personal data is processed (s. Sec. 2.5). This can easily be the case if massive amounts of textual data are crawled for general pre-training and, e.g., e-mail addresses are unintentionally stored in the process. Regarding fine-tuning, images taken inside a company's premise, e.g., for an action recognition task, may contain employees' faces or specific body characteristics. Although most data points will typically represent non-personal data, the GDPR applies to the whole dataset as long it is not possible to make a clear distinction between both and process them separately.

The legal basis (Art. 6 (1)) for data processing in the context of the training of VLFMs may, in particular, be consent (lit. a) or the necessity to process personal data ”for the purposes of the legitimate interests pursued by the [data] controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject” (lit. f). While getting consent is not practical in the case of personal data gathered via web crawling, it may constitute a viable option regarding internal data. However, if employees' personal data is going to be used, the workers‘ representatives typically negotiate strict rules on its processing. Additionally, the possibility of withdrawing consent at any time (Art. 7 (3)) creates considerable legal uncertainty. Therefore, relying on (lit. f) seems advisable in both cases. Accordingly, the processing must be *necessary* to serve a *legitimate interest*, and the interests or fundamental rights and freedoms of the data subjects must not *outweigh* the legitimate interests of the controller or a third party on a case-by-case basis.

The *legitimate interests* of the controller can include non-commercial and commercial interests. If the pre-trained VLFM is made available to the general public, as it is often the case with FMs, an interest in the information and education of the general public may constitute an additional legitimate interest. The *necessity* of data processing presupposes that legitimate interests cannot be achieved with less intrusive measures. In addition to questions of possible anonymization of training data, the processing of personal data must thus be reduced to the extent necessary for effective training. Factors to be considered in the case-based *balancing of interests* include the type and scope of processing, the type of data, the degree of transparency vis-à-vis the data subjects or the implementation of technical and organizational measures by the controller to uphold the principles laid out in Art. 5 (1). Here, it must be kept in mind that data collection via crawling, e.g., vast amounts of non-structured internal documents or the web, is neither transparent (lit. a) nor does including such data in a potentially distributed dataset preserve confidentiality (lit. f). Furthermore, the training of VLFMs, which are initially open in terms of their application scope, may also conflict with the principle of data collection for specified and explicit purposes (lit. b). Since those points relate to all FMs, the advice given by legal scholars in the broader context of FMs is relevant. They put particular emphasis on a robust data governance process, which should also take into account the requirements for training, validation, and test data of high-risk AI systems, as laid out in the AI Act (Art. 10 (2)), if applicable [14]. How such a process should look like in practice to reliably outweigh the data subjects' rights remains, however, an open question.

3.1.2. Directive on Copyright in the Digital Single Market

Since copyright law is determined by the legislation of the EU's member states, regulations diverge. Nevertheless, a shared key objective is to ensure that authors receive appropriate compensation for the commercial use of their work. This requires provisions that determine the eligibility of a work for protection, its attribution to an author, and rights that protect the author against unauthorized exploitation. The lawful use of copyrighted work is typically achieved via authorization, often in exchange for money. However, statutory exceptions may permit the use of copyrighted work even without the author's authorization. For the training of VLFMs, an exception for data mining is of interest. Conveniently, such an exception is the subject of the Directive on Copyright in the Digital Single Market and, therefore, must be adopted by all member states.

Art. 2 (2) defines "text and data mining" as "any automated analytical technique aimed at analyzing text and data in digital form in order to generate information which includes but is not limited to patterns, trends, and correlations". Art. 4 (3) states that "[t]he exception [...] shall apply on condition that the use of works [...] has not been expressly reserved by their rightholders in an appropriate manner, such as machine-readable means in the case of content made publicly available online". Art. 3 provides a similar exception for "research organisations and cultural heritage institutions" (defined in Art. 2 (1, 3)), which, notably, does not contain a condition like the one of Art. 4 (3). Broadly speaking, it can thus be stated that text and data mining is generally permitted for commercial purposes as long as the copyright holder does not opt-out. In the case of collection for scientific research, opt-outs are not possible.

The EU's first major court case dealing with data mining for the purpose of AI training recently took place in Germany. A photographer filed a lawsuit against a nonprofit organization focused on advancing AI research by providing open datasets for model training, including billions of image-text pairs. Among the images was one that belonged to the claimant. He had used a website for distribution which explicitly prohibited the use of "automated programs". The court had to decide whether the exceptions for text and data mining, adopted in § 44b (→ Art. 4) and § 60d (→ Art. 3) of the German 'Urheberrechtsgesetz' (UrhG), protected the nonprofit of copyright infringement. It denied a copyright infringement based on § 60d UrhG (→ Art. 3), as it deemed the compilation of a dataset for the subsequent generation of new knowledge to classify as scientific research. More interestingly, however, the court advocated for a broad interpretation of "machine-readable" (§ 44 (3) → Art. 4 (3)). It hinted that, depending on the circumstances of the case and the technical development at the time, opt-outs formulated in natural language and placed on a website can be sufficient.

As national laws based on the directive are supposed to be interpreted uniformly within the EU, the ruling is of great importance. Nevertheless, as the German court did not refer any of the issues raised to the ECJ, other courts may diverge from its assessments. Regarding the collection of data for the training of VLFMs in Europe by commercial providers who cannot rely on the broader exception based on Art. 3, it suggests that crawling in compliance with "robots.txt" files can become problematic. Rather, the text of the web pages should be checked for exclusions using NLP methods currently deemed SOTA. This significantly slows down the collection of training data.

3.2. Problem: Pre-Trained Models

We have stated that pre-training is problematic from a legal point of view because the reliance on heavy web crawling can easily lead to the inclusion of protected works and/or personal data. Consequently, one could assume that responsibility for data protection and copyright law can be eliminated if models are not pre-trained in-house, but existing models are only fine-tuned on smaller, task-specific datasets. However, even if those datasets are carefully compiled and do not include protected works nor personal data, the sole usage of a model that was trained on problematic data might result in legal responsibilities. Since overparametrized neural networks like VLFMs are able to reproduce training data (which was highlighted by training data extraction attacks, see [15]) and must therefore, albeit often in compressed form, store it. Legal scholars and data protection authorities have argued that both copyright law [16] and the GDPR [17] have to be considered when pre-trained models are distributed, copied, or used. In both legal areas, these assessments are contested [18, 19]. Given the absence of court rulings, this, nevertheless, adds legal uncertainty even if previously pre-trained (or fully fine-tuned) models are employed.

3.3. Practical Application of VLFMs

When fine-tuned VLFMs are employed for practical applications, providers and deployers should pay attention to the AI Act. Since VLFMs can be adapted towards various downstream tasks, the individual risk level must be determined in light of the use case (s. Sec. 3.3.1). While certain systems are generally prohibited (Art. 5), high-risk AI forms the centerpiece of the regulation (s. Sec. 3.3.2). Since VLFMs may be used to interact with workers directly and can generate artificial content, they are also subject to the requirements of Art. 50 (s. Sec. 2.4).

3.3.1. Risk Level Assessment

As per to Art. 6 (1), a high-risk AI system exists if two conditions are met. Firstly, the AI system must be used as a safety component of a product listed in Annex I or be such a product itself. Secondly, the relevant EU product safety regulations must stipulate that a third-party conformity assessment must be carried out for the product. In the scope of tangible manufacturing steps, vision-based safety systems are of particular concern, rather than VLFMs assisting in in-tangible processes, e.g., design engineering. The second category of high-risk AI systems (Art. 6 (2)) refers to stand-alone AI systems that pose a high risk to health, safety, or fundamental rights. If AI systems are intended to be used in an area listed in Annex III of the AI Act, which includes critical infrastructure, they are also considered to be high-risk.

3.3.2. High-Risk AI

The core safety requirements of the AI Act are imposed on high-risk AI. They include risk management, data governance, technical documentation, record-keeping, transparency, human oversight, as well as general accuracy, robustness, and cybersecurity (Art. 8–15). As per Art. 16, these primarily address the *providers*, who are to ensure compliance with the requirements by means of a quality management system, documentation, and logs (Art. 16-19). However, the AI Act also obligates importers, distributors, and deployers. As per Art. 26, *deployers* must ensure that the systems are used as intended and with sufficient human supervision, that input data is relevant and sufficiently representative given the purpose, and that logs are kept. Very relevant for industrial applications is Art. 26 (6): before using high-risk AI at the workplace, "deployers who are employers shall inform workers' representatives and the affected workers that they will be subject to the use of the high-risk AI system."

For the enforcement of the requirements, the AI Act relies on a decentralized approach. The *providers* generally subject their systems to a conformity assessment based on internal control in accordance with Art. 43 (1) (lit. a) in conjunction with Annex VI. For VLFMs in industrial applications, the internal control and documentation obligations may be included into pre-existing quality management systems. Optionally, providers can also choose the external conformity assessment procedure with the involvement of a notified body described in Art. 43 (1) (lit. b) in conjunction with Annex VII of the AI Act. The latter is currently only mandatory for systems operating with biometric data (Annex III No. 1). *Deployers* establish and document a post-market monitoring system (Art. 72) and, according to Art. 26 (5), inform the provider/distributor and surveillance authority when necessary.

4. Conclusion and Outlook

The EU's legal framework imposes relatively strict copyright and data protection guardrails for the training of VLFMs. Web crawling, which becomes necessary if models are trained "from scratch," is especially problematic. However, fine-tuning may also raise concerns regarding the GDPR. Greater legal certainty is urgently needed and will hopefully be provided by upcoming court decisions. As of now, developers should use NLP to scan for opt-outs during crawling (1) and limit the use of personal data via data governance policies (2).

Since VLFMs are GPAI models, the requirements laid out in Art. 51–56 of the AI Act have to be taken into account. Given that fined-tuned models will often be intended to interact with humans and can generate content, the transparency requirements of Art. 50 (1, 2) are relevant as well. Depending on the specific use case, VLFMs may constitute high-risk AI. If so, this results in extensive safety requirements (Art. 8–15).

Larger corporations often possess organizational frameworks to ensure compliance with various regulations. These existing structures can be extended to monitor and enforce the regulatory compliance for high-risk AI. In contrast, small and medium-sized enterprises are more reliant on clear standards and guidelines that set out concrete compliance measures for the practical use of AI. These are currently lacking.

References

[1] R. Bommasani, D. A. Hudson, E. Adeli, et al., On the Opportunities and Risks of Foundation Models (2021). doi:10.48550/arxiv.2108.07258.

[2] H. Golnabi, A. Asadpour, Design and application of industrial machine vision systems, Robotics and Computer-Integrated Manufacturing 23 (6) (2007) 630–637. doi:10.1016/j.rcim.2007.02.005.

[3] K. Moenck, A. Wendt, P. Prünte, et al., Industrial Segment Anything – a Case Study in Aircraft Manufacturing, Intralogistics, Maintenance, Repair, and Overhaul (Jul. 2023). doi:10.48550/arxiv.2307.12674.

[4] OpenAI, GPT-4V(ision) System Card (Sep. 2023). URL https://cdn.openai.com/papers/GPTV_System_Card.pdf

[5] A. Dubey, A. Jauhri, A. Pandey, et al., The Llama 3 Herd of Models (Aug. 2024). doi:10.48550/arxiv.2407.21783.

[6] F. Zeng, W. Gan, Y. Wang, et al., Large Language Models for Robotics: A Survey (Nov. 2023). doi:10.48550/arxiv.2311.07226.

[7] Y. Tian, A. Liu, Y. Dai, et al., Systematic synthesis of design prompts for large language models in conceptual design, CIRP Annals 73 (1) (2024) 85–88. doi:10.1016/j.cirp.2024.04.062.

[8] J. Koch, D. Jevremovic, K. Moenck, T. Schüppstuhl, A digital assistance system leveraging vision foundation models & 3D localization for reproducible defect segmentation in visual inspection, Procedia CIRP 130 (2024) 387–397. doi:10.1016/j.procir.2024.10.105.

[9] C. Picard, K. M. Edwards, A. C. Doris, et al., From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design (Aug. 2024). doi:10.48550/arxiv.2311.12668.

[10] K. Moenck, D. T. Thieu, J. Koch, T. Schüppstuhl, Industrial Language-Image Dataset (ILID): Adapting Vision Foundation Models for Industrial Settings, Procedia CIRP 130 (2024) 250–263. doi:10.1016/j.procir.2024.10.084.

[11] Regulation (EU) 2024/1689 (Jun. 2024). URL http://data.europa.eu/eli/reg/2024/1689/oj

[12] Regulation (EU) 2016/679 (May 2016). URL http://data.europa.eu/eli/reg/2016/679/2016-05-04

[13] Directive (EU) 2019/790 (May 2019). URL http://data.europa.eu/eli/dir/2019/790/oj

[14] J. Hüger, Die Rechtmäßigkeit von Datenverarbeitungen im Lebenszyklus von KI-Systemen (2024).

[15] N. Carlini, F. Tramer, E. Wallace, et al., Extracting Training Data from Large Language Models (Jun. 2021). doi:10.48550/arxiv.2012.07805.

[16] T. W. Dornis, S. Stober, Urheberrecht und Training generativer KI-Modelle, Nomos Verlagsgesellschaft mbH & Co. KG, Baden-Baden, 2024. doi:10.5771/9783748949558.

[17] M. Hansen, B. Walczak, Die KI zaubert nicht. Warum ein Personenbezug in LLMs erhalten bleibt (2024).

[18] L. Käde, Kreative Maschinen und Urheberrecht, Nomos Verlagsgesellschaft mbH & Co. KG, Baden-Baden, 2021.

[19] Hamburgische Beauftragte für Datenschutz und Informationsfreiheit (HmbBfDI), Diskussionspapier: Large Language Models und personenbezogene Daten (2024).