AI GOVERNANCE GUIDE

International Federation of Reproduction Rights Organisations

CONTENTS

ABOUT IFRRO'S AI GOVERNANCE GUIDE

IFRRO is an independent, non-profit membership association that facilitates, at international level, the collective management of reproduction and other rights in text- and image-based works through cooperation among its member Reproduction Rights Organizations (RROs). IFRRO brings together over 160 members from some 90 countries worldwide, representing many millions of authors, visual artists and publishers of books, journals, newspapers, magazines and printed music. IFRRO's mission is to develop and promote effective collective rights management so as to ensure that the copyrights of authors and publishers are properly valued through the lawful and remunerated use of text- and image-based works. In pursuing this mission, IFRRO is guided by the values of trust, diversity and commitment, and acts with tolerance and respect in its engagement with members, policymakers and stakeholders.

This AI Governance Guide is the result of a close collaboration between IFRRO and Professor Dr. Eleonora Rosati, a leading expert on copyright and AI, building on her longstanding collaboration with IFRRO on the AI Advocacy Toolkit. It reflects IFRRO's experience supporting lawful access to content through efficient licensing solutions, and its conviction that copyright is not an "add-on" to AI governance, but a core pillar of responsible and sustainable AI development and deployment. High-quality content is created by human beings—authors, visual artists, journalists, researchers, educators and publishers—and this content underpins the knowledge, culture and innovation that AI systems depend on. If AI ecosystems can mobilize major investments in infrastructure, computing power and energy, they must also recognize and value the essential role of copyright-protected works as a primary input into AI systems.

Over the past year, AI governance has moved rapidly to the top of national, regional and international policy agendas, including through strategic plans, white papers and high-level reports. Yet, across many of these initiatives, the copyright dimension remains insufficiently developed, despite its direct relevance to the full AI lifecycle from the acquisition and use of training material, to fine-tuning, deployment, downstream uses, and the circulation of outputs in the marketplace. This gap creates legal uncertainty, weakens trust, and risks undermining the creative and publishing ecosystems that are necessary for the continued availability of reliable, diverse and high-quality content.

Against this backdrop, this Guide aims to place copyright at the center of AI governance discussions and to support policymakers, developers, deployers, rightsholders and other stakeholders in adopting workable, future-oriented solutions. It advances a balanced approach that is both pro-copyright and pro-innovation: enabling the development of powerful and reliable AI systems while ensuring that the use of protected works is lawful, transparent and appropriately remunerated.

Further information on IFRRO's work on artificial intelligence, including member-related actions, initiatives and collective licensing solutions, is available via IFRRO's dedicated AI page: https://ifrro.org/page/Artificial-Intelligence-public/

TABLE OF ABBREVIATIONS

AI	Artificial Intelligence
AI Act	Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024 (EU)
CDPA	Copyright, Designs and Patents Act 1988 (UK)
CJEU	Court of Justice of the European Union
DSM Directive	Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, OJ L 130, 17.5.2019 (EU)
GPAI	General-Purpose AI
InfoSoc Directive	Directive 2001/29/EC of the European Parliament and of the Council of 22 May 2001 on the harmonisation of certain aspects of copyright and related rights in the information society, OJ L 167, 22.6.2001, pp. 10–19 (EU)
TDM	Text and data mining

EXECUTIVE SUMMARY

The rapid – or rather: explosive – rise of generative AI has been affecting several aspects of societies, economies, and business models around the world. Following the release of the IFRRO AI Toolkit, this Guide intends to: (i) explore the interplay between AI development and copyright law and persona and style protection; (ii) identify legislative, judicial and policy responses and emerging governance models of core aspects thereof in selected jurisdictions around the world; and (iii) formulate recommendations for the governance and sustainable progression of AI-driven innovation.

Parts I and II discuss key aspects relating to the transition from AI training to output generation. Specifically, Part I discusses the border between licensing, infringement and exceptions to copyright and related rights in the context of AI training (Section A), transparency obligations relating to AI training materials (Section B), and the issues that the increased ease with which synthetic media and deepfakes may be generated and circulated raise vis-à-vis persona protection, including in the publishing context (Section C). Part II focuses on the protectability of AI-generated and -assisted outputs under copyright (Section A) and labelling requirements (Section B). Finally, Part III identifies governance responses to the aspects analysed in Parts I and II and provides key recommendations that seek to ensure that AI-based innovations can flourish in a way that is respectful of rights granted under copyright, related rights, and the protections afforded to one's persona, so as also to comply with the mandate – found in inter alia fundamental and human rights instruments (including the European Convention on Human Rights and the Charter of Fundamental Rights of the European Union) and international copyright instruments (including the WIPO Copyright Treaty) alike – that a fair balance of rights and interests is ultimately struck.

Overall, while the advent and implementation of generative AI raises similar issues worldwide, the provision of legislative and policy responses has progressed at a different pace, with some jurisdictions having already adopted governance models and others (the majority) having yet to take specific initiatives. All this leaves future regulation and governance of AI undefined – yet bound by pre-existing legal concepts – in most cases. In light of this background, the main findings and recommendations are the following:

I. From AI training to output generation

- Training AI models requires access to and use of large amounts of works and other subject-matter often protected by copyright and other rights. The exclusive rights granted by copyright are preventive in nature: in turn, lacking an applicable exception or defence, the undertaking of acts restricted by such rights is unlawful.

- There is no exception or open-ended fair use-style defence that covers the entirety of AI training acts in each and every instance, including but not limited to commercial contexts. As a result, a licensing approach is preferable and practical. The latter is also supported by the existence of both direct and collective licensing solutions.

- AI developers must be transparent about works and other subject-matter used in training, having secured the necessary permissions or licenses where needed. Legislators should mandate robust transparency requirements to help rightholders enforce their rights and prevent unauthorised uses of their protected material.

- Synthetic media produced by generative AI presents challenges, as international and regional frameworks regarding outputs are thin. Some jurisdictions impose labelling requirements and more are considering them. At the national level, existing or even new rights may be used to prevent unauthorised uses of one's likeness and other personal attributes, which is also is key to serve democracy goals and reduce the risks of manipulation and misinformation.

- In particular cases, where they are sufficiently distinctive and personal, styles can qualify for protection as long as they may be regarded as original works. Hence, copying one's style might be also actionable under copyright law, as well as other substantive rights and general causes of action.

- Authors and rightholders must be able to enforce their rights both when their works are used without permission for training purposes and when AI outputs unduly imitate works and styles.

II. AI-generated outputs

- Most jurisdictions require human authorship for copyright protection and AI alone cannot be deemed an author. The extent of human contribution determines whether a work qualifies for copyright protection. Only human-created elements within AI-generated and -assisted works can be protected, also emphasising the need to safeguard human creativity.

- Transparency and labelling obligations for AI-generated content are being implemented to foster user trust and recognise the value of human works. Legislators should introduce labelling requirements and ensure the effectiveness thereof. Labelling practices adopted at the level of individual companies or organisations are helpful and should be also encouraged.I. From AI training to output generation

I. FROM AI TRAINING TO OUTPUT GENERATION

This Part focuses on some key issues facing AI training from a copyright and related rights perspective. It also discusses the problems raised by unauthorised synthetic media and deepfakes, that is: AI-generated and -assisted outputs that reproduce one's likeness and other personal attributes. First, the border between licensing and exceptions / open-ended defences (fair use) to copyright and related rights in the context of AI training is discussed (Section A). Then, it is the turn of transparency obligations in AI training processes and data (Section B). Finally, the issues that the increased ease with which synthetic media may be generated and circulated raise vis-à-vis persona protection, including style imitation, are examined (Section C).

A. Licensing vs exceptions for AI training

AI training requires access to and use of large quantities of works and subject-matter, often called 'data' and more often than not protected by inter alia copyright, related rights, and sui generis rights. In this context, studies also show that human-created works, as opposed to synthetic content, is better for high-quality AI results (see, e.g., I Shumailov and Others, 'AI models collapse when trained on recursively generated data' (24 July 2024) Nature).

Against the background above, much of the global debate regarding AI training centres on how to balance rightholders' control, on the one hand, with the use of copyright works and other protected-subject matter by AI developers without permission, on the other. In other words: licensing versus exceptions.

1. TDM exceptions and open-ended defences (fair use), also in light of the three-step test/fair dealing requirements

As discussed in greater detail in the IFRRO AI Toolkit, some jurisdictions around the world have adopted exceptions allowing TDM under certain conditions:

- **EU**: Articles 3 and 4 of the DSM Directive provides for two distinct exceptions that only cover acts of extraction and reproduction for TDM purposes. The former is for the benefit of research organisations and cultural heritage institutions, while the latter does not have restrictions on beneficiaries and does not in itself exclude commercial TDM. That said, Article 4(3) of the DSM Directive allows rightholders to reserve their rights and, thus, opt out of the application of the exception. The modalities of such rights reservation are currently a matter of discussion and controversy, including in the aftermath of the release of the General-Purpose Code of Practice. The few cases decided so far in, respectively, Germany, Netherlands, Hungary, and Denmark have failed to settle this issue, which therefore remains contentious:

    - **Germany and Denmark**: At first instance, in Knesche v. LAION, 310 O 227/23 (2024), the District Court of Hamburg noted that a rights reservation done using natural language shall be deemed valid and enforceable (M Brüß, 'German Court finds LAION's copying of images non-infringing' (28 September 2024) The IPKat). The Munich Regional Court shared the same conclusion in the more recent judgment in GEMA v. OpenAI, 42 O 14139/24 (2025). Similarly, in BoligPortal v. ReData, BS-42485/2025-SHR (2025) the Danish Maritime and Commercial Court adopted the view that the claimant's clear prohibition against data mining and privacy policy constituted a valid TDM reservation.

    - **Germany, Netherlands and Hungary**: In DPG MEDIA BV and Others v KNOWLEDGE EXCHANGE BV, C/13/737170 / HA ZA 23-690, NL:RBAMS:2024:6563 (2024), the Amsterdam District Court found that an appropriate reservation would be one done in the programming language robots.txt (A Cerri, 'Dutch court holds that TDM opt-out must be done by "machine-readable" means' (16 February 2025) The IPKat). The same approach was endorsed by the Hungarian Municipal Court of Appeals in Case 9.Pf.20.353/2024/6-II (2024) (P Mezei, 'Third European court decision on the general purpose TDM exception is out' (8 May 2025) Kluwer Copyright Blog). In the 2025 appellate decision in Knesche v. LAION, 5 U 104/24, contrary to the first instance court, the OLG Hamburg held that a reservation done through natural language employed in the terms of service of a website does not satisfy the requirement of machine readability.

All the above said, the consideration that machines understand natural language suggests that a **rights reservation done in natural language** is sufficient (R Kaufman, 'AI rights reservation: Human readable is machine readable – An interview with Haralambos ("Babis") Marmanis' (17 February 2025) The Scholarly Kitchen).

Like Company v. Google CJEU referral: At the time of writing, a referral for a preliminary ruling (Like Company v Google, C-250/25) is pending before the CJEU. The questions referred relate to the interplay between copyright and AI. Among other things, the Budapest District Court (Hungary) is asking if AI training engages the right of reproduction and, if so, whether the TDM exception in Article 4 of the DSM Directive applies. As it is discussed elsewhere (E Rosati, 'CJEU receives first referral on chatbots and copyright' (26 May 2025) The IPKat):

- Whether AI training engages the right of reproduction: This question appears to revisit the vexed issue whether training is even relevant under copyright law. At least in Europe, an answer in the affirmative has been already given through the adoption of specific TDM exceptions to inter alia the right of reproduction. The CJEU should answer this question accordingly.

- Whether the TDM exceptions apply to AI training: While some commentators have also recently opined that the TDM exceptions were not designed with AI training in mind (see, e.g., N Lucchi, Generative AI & Copyright: Balancing Creative Rights, Legal Integrity, and Accountability in the AI Age (4 June 2025)), the adoption of inter alia the AI Act appears to posit that the TDM exceptions *may* be relevant in the context of AI training. Nevertheless, as it has been explained, there is no exception under copyright and other rights that encompasses 'AI training' as such. In the EU, there are exceptions to specified restricted acts for specified purposes relating to TDM, which are all subject to several different conditions, including (i) lawful access to the protected content in question and (ii) the requirements of the three-step test. The CJEU should uphold the specific and limited scope of TDM exceptions.

UK: A non-commercial text and data analysis exception (section 29A) was introduced into the CDPA in 2014 by relying on Article 5(3)(a) of the InfoSoc Directive. A potential legal reform is currently being considered.

In late 2024, the UK government launched its Copyright and AI consultation, seeking views on the application of UK copyright law to the training of AI models. The proposed reformed UK approach could be closely modelled after EU law: a new exception might be introduced into the CDPA to allow TDM on lawfully accessed content for any purpose, subject to the possibility for rightholders to reserve their rights. Greater transparency about the sources of training material might be also imposed, again taking inspiration from EU law (notably Article 53(1)(d) of the AI Act). The consultation closed on 25 February 2025, attracting over 11,500 responses and the input of several stakeholders, including high-profile artists opposing reform plans (L Kuenssberg, 'Elton brands government 'losers' over AI copyright plans' (18 May 2025) BBC; Creative Titans and Industry Leaders Urge the Prime Minister to Support UK Creativity and Economic Growth by Enforcing Copyright Law (10 May 2025), M Sellman, 'Abba's Bjorn Ulvaeus joins chorus against AI copyright overhaul' (29 April 2025) The Times. At the time of writing, the official response of the UK Government and any

follow-up initiative remain pending (Department for Business and Trade, The UK's Modern Industrial Strategy 2025 (23 June 2025), p. 126). Despite vigorous stakeholder advocacy, the Data (Use and Access) Act 2025 does not contain any transparency obligations, the justification being that this legislation differs and is separate from an official response to the already mentioned public consultation.

In late 2025, the High Court of England and Wales delivered its judgment in Getty Images (US) Inc & Ors v Stability AI Ltd (Rev1) [2025] EWHC 2863 (Ch) (04 November 2025). The judgment does not concern Stability AI's potential primary liability for copyright infringement, and does not interpret or apply section 29A CDPA, which thus remains untested at the time of writing. The decision has been appealed and the case is in progress at the time of writing. Then, in December, it was revealed that an overwhelming majority (95%) of respondents to the 2024-25 public consultation was against loosening copyright law to facilitate unlicensed AI training. On 15 December 2025, the UK secretary of state for science, innovation and technology told parliament that there was "no clear consensus" and the government would "take the time to get this right" (R Booth, 'Boost for artists in AI copyright battle as only 3% back UK active opt-out plan' (16 December 2025) The Guardian).

In March 2026, the House of Lords' Communication and Digital Committee published its report on AI, Copyright and the Creative Industries, which explores how the UK can enable responsible, licensing-based AI development that protects UK creators while supporting innovation. Among the main findings, there is that a commercial TDM exception with an 'opt-out' mechanism would not deliver meaningful control for rightholders and would remove incentives for developers to license content for AI training. Creating the conditions for a thriving licensing market for AI training is recommended instead. This can be achieved by supporting the development and adoption of technical standards and tools that could overcome barriers to licensing at scale.

Asia: TDM exceptions exist in the copyright laws of, e.g., **Japan** (Article 30-4 of the Copyright Act) and Singapore (Section 244 of the Copyright Act 2021). **Singapore** also envisages an open-ended fair use defence closely modelled upon, although not only mirroring, the U.S. doctrine (section 191 of the Copyright Act).

Insofar as **South Korea** is concerned, in 2011 a fair use provision (currently Article 35-5 of the Copyright Act, as last amended in 2023) was adopted, also modelled upon the corresponding U.S. fair use doctrine, yet with some crucial differences from U.S. law, including the express reference in the provision (para 1) to the requirements of the three-step test. In A Guide on Generative AI and Copyright, released by the Ministry of Culture, Sports and Tourism and the Korea Copyright Commission in 2024, the example of web scraping for the purpose of creating datasets to be subsequently employed in AI training processes is specifically discussed, the conclusion being that such an activity might be infringing under copyright law if undertaken without a licence (p. 10). The Guide further highlights how the applicability of the national fair use doctrine to unlicensed AI training remains a matter of contention among commentators. Because of this, it recommends "to prevent possible disputes by securing authorization from the rights holder prior to use" (p. 17).

None of the exceptions adopted in the jurisdictions above (i) are for 'AI training' as such, nor do they, in any case, (ii) extend to output generation.

In turn, if an AI-generated or -assisted outputs reproduces training data, questions of potential liability shall arise. These are discussed in greater detail in the IFRRO Toolkit (see also E Rosati, 'Infringing AI: Liability for AI-generated outputs under international, EU, and UK copyright law' (2025) 16(2) European Journal of Risk Regulation 603; regarding specifically the case of Singapore, see Artificial Intelligence Output and Copyright Infringement in Singapore, noting that: "AI tools can create stories, images, songs and more. But what happens if such output looks too much like someone else's original work? As a user, developer or deployer of AI systems, you could be liable for copyright infringement by using AI irresponsibly.") As regards the USA, in a recent case, the U.S. District Court for the Southern District of New York denied the defendant's motion to dismiss claims that certain AI-generated outputs infringe copyright (Case 1:23-cv-08292-SHS-OTW).

In the vast majority of countries, no TDM exceptions have been adopted, including in China. In countries where such reforms have been discussed, either there is an impasse or eventually a decision has been reached to abandon reform plans as detrimental to its creative industries (see also J Gong, 'What you need to know about China's new Generative AI measures' (3 August 2023) Bird & Bird):

- An example of the former is **Hong Kong**, where in 2022 the Government announced that, because of the diversity of views among stakeholders, "rushing into incorporating these issues in the amendment bill" would not be recommended (Legislative Council Panel on Commerce and Industry, Updating Hong Kong's Copyright Regime – Outcomes of Public Consultation and Proposed Way Forward (19 April 2022), p. 19. In 2024, a new public consultation was nevertheless launched, reviving the discussion around introducing a new TDM exception (Commerce and Economic Development Bureau Intellectual Property Department, Copyright and Artificial Intelligence – Public Consultation Paper (2024), p. 30 ff). At the time of writing no further developments have occurred.

- An example of the latter is **Australia**, where the Government announced that no such exception would be introduced into Australian law (Albanese Government to Ensure Australia is Prepared for Future Copyright Challenges Emerging from AI (26 October 2025)).

In the **USA**, no legislative reform has been undertaken. Ongoing litigation around unlicensed AI training is thus testing the boundaries of the fair use doctrine:

- In the first case decided on unlicensed AI training, Thomson Reuters v. Ross Intelligence, No. 1:20-CV-613-SB (D. Del. Feb. 11, 2025), the U.S. District Court for the District of Delaware concluded that the defendant's conduct did not qualify as fair use. Ross Intelligence's training of a competing AI model to Thomson Reuters' Westlaw infringed copyright in the latter's headnotes. At the time of writing, Ross Intelligence's appeal is pending.

- In the more recent order issued in Richard Kadrey, et al. v. Meta Platforms, Inc., No. 23-cv-03417-VC (N. Cal. Jun. 25, 2025), Judge Chhabria granted summary judgment to Meta on the plaintiffs' claim that the company had infringed copyright by training its models with their books on consideration that the plaintiffs had made the wrong argument and failed to develop a record in support of the right case. As a result, not only – "in the grand scheme of things" – the consequences of the ruling are limited, but what is particularly relevant is the point, made by the judge, that "in most cases" the unlicensed use of protected content for training purposes would be regarded as illegal, also considering that, if these models are expected to generate billions, even trillions, of dollars for their developers, it is "ridiculous" to claim that such companies cannot "figure out a way to compensate copyright holders" for the use of their content, if this is required for training purposes. Judge Chhabria also criticised the earlier order issued by Judge Alsup in Andrea Bartz, et al., v. Anthropic, No. C 24-05417 WHA (N. Cal. Jun. 23, 2025), especially that he had appeared to overlook the substantial displacement effects of AI-generated content flooding the market and the reduction of incentives to create in the future. In 2025 Anthropic decided to settle its litigation with Bartz and Others for USD 1.5 billion, this being the largest settlement in the history of U.S. copyright litigation (D Hansen, 'The Anthropic settlement – what it is and isn't (and who could get paid)' (7 September 2025) Authors' Alliance).

- The displacement effect of AI-generated content vis-à-vis human-created work – and thus its relevance under the fourth fair use factor – has been also highlighted in literature, including empirical and qualitative research (T Chakrabarty – JC Ginsburg – P Dhillon, 'Readers prefer outputs of AI trained on copyrighted books over expert human writers' (last revised 9 November 2025) Columbia Public Law Research Paper No. 5606570).

No copyright exception – including those permitting TDM – covers the entirety of AI development in each and every case (E Rosati, 'Copyright exceptions and fair use defences for AI training done for "research" and "learning," or the inescapable licensing horizon' (2025) European Journal of Risk Regulation). All existing TDM exceptions only encompass specified restricted acts. AI training – by default – requires the making of copies at multiple stages, with the result that the right of reproduction is engaged several times. Other exclusive rights are also likely to be undertaken in the context of training processes.

The exclusive right of reproduction encompasses reproduction "in any form". It is irrelevant if the copies made are permanent or merely temporary. If a copy is made at some point in the process of developing an AI model, that copy is the result of a copyright-relevant act of reproduction. Hence, those that have argued that the process of development of AI models does not involve copyright-relevant reproductions – either because the copies are 'intermediate' copies, or that the AI models are merely being 'inspired' by existing works and subject-matter, or even referring to the idea/expression dichotomy and arguing that copying only involves 'non-expressive' parts of works – are incorrect, both from a technical perspective and the perspective of copyright law. The diagram below offers a simplified guide of the stages of AI model development that could involve copyright-relevant acts:

Scraping/crawling of protected content online

- Right of extraction (if content included in database)
- Right of reproduction

Online publication of dataset

- Right of reproduction
- Right of communication/ making available to the public

AI training

- Right of extraction (if content included in database)
- Right of reproduction

Creation of dataset

- Right of reproduction

Downloading of dataset

- Right of reproduction

Use of resulting model

- Right of reproduction
- Right of adaptation
- Right of communication/ making available to the public
- Right of distribution

Like all exceptions, the TDM ones also require (1) lawful access to the protected content and (2) compliance with the three-step test and, having regard to the UK, fair dealing:

- **Lawful access** means authorised release of the protected work or subject-matter by the relevant rightholder before it can be used by others without permission. It is a general requirement under copyright law, regardless of whether a specific mention in this sense is provided in the statutes.
  - **EU:** The CJEU has consistently held that the availability of exceptions in general presupposes prior lawful access, further linking such a requirement to the three-step test (e.g., CJEU, ACI Adam, C-435/10, para 38).
  - **USA:** The Copyright Office has concluded that training of AI models on content that is known to be unlawful weighs against a finding of fair use (U.S. Copyright Office, Copyright and Artificial Intelligence – Part 3: Generative AI Training (Pre-publication Version) (May 2025), p. 52). Overall, under U.S. law, "downloading and storing pirated content is clearly not legal, even if it is to be used for a fair use purpose." (H Stephens, 'When the end does not justify the means, Anthropic's $1.5 billion lesson' (13 September 2025) Hugh Stephens Blog).
  - **Japan:** The same position has been indicated to apply under Japanese law (Copyright Division, Agency for Cultural Affairs, Japan), "General Understanding on AI and Copyright in Japan" -Overview- (published by the Legal Subcommittee under the Copyright Subdivision of the Cultural Council) (May 2024), p. 11).
- The **three-step** test consists of three cumulative requirements – (i) certain special cases, (ii) no conflict with normal exploitation, (iii) no undue prejudice to rightholders' legitimate interests – which need to be assessed in accordance with their logical order, that is as steps. Before determining whether a certain exception unreasonably prejudices the legitimate interests of the concerned rightholder, it is necessary to consider whether there is a conflict with a normal exploitation of protected content and, prior to that, whether the exception at hand only applies to certain special cases (WTO, United States – Section 110(5), §6.160; see also Guide to Copyright and Related Rights Treaties Administered by WIPO and Glossary of Copyright and Related Rights Terms, para 85). As AI licensing becomes more prevalent, the 'normal exploitation' of a copyright work or other protected subject-matter is indeed the one controlled by the relevant rightholder.

Applying all of the above to AI training, **not only is there no TDM exception (or open-ended fair use doctrine) covering in every instance the entirety of it, but such an exception also would not be possible considering the requirements of the three-step test / fair dealing** (E Rosati, 'TDM exceptions (not just the three-step test) don't allow all unlicensed AI development' (13 October 2025) The IPKat). A different interpretation would (i) arguably render an exception applicable well beyond 'certain special cases' – thus disqualifying it at the outset given that, as noted, the steps of the three-step test must be considered in sequence – and (ii) be at odds with the second and third step.

2. The importance of licensing

Despite TDM exceptions and arguments that open-ended fair use-style doctrines are generally applicable to unlicensed AI training (see, e.g., MW Carroll, 'Copyright and the progress of science: Why text and data mining is lawful' (2019) 59 UC Davis Law Rev 893; MA Lemley – B Casey, 'Fair learning' (2021) 99 Tex L Rev 743), no exception / open ended, fair use-style doctrine fully covers the development and deployment of AI models in each and every instance. For rightholders and AI developers, a licensing approach is thus both practical and necessary as it is through licensing that concerned parties are provided certainty regarding their rights and responsibilities with respect to the use of protected works and subject-matter for training.

A licensing market for AI training already exists and continues to evolve. As reported in the media, agreements have been already concluded between rightholders and AI developers (see Generative AI Licensing Agreement Tracker, Ithaka S+R; the AI Licensing Deals as listed here; AI Licensing by AI Companies). Collective management organisations have also been developing and launching collective licences for GenAI. Licences have been, for example, unveiled by CCC, CLA, and STIM.

Going forward, the applicability of exceptions could (and should) be likely confined, as a matter of practice, to very circumscribed phases of AI development. **Licensing, not exceptions, is the most practical and – above all – practicable framework for accommodating AI development and striking the required 'fair balance' with copyright protection.**

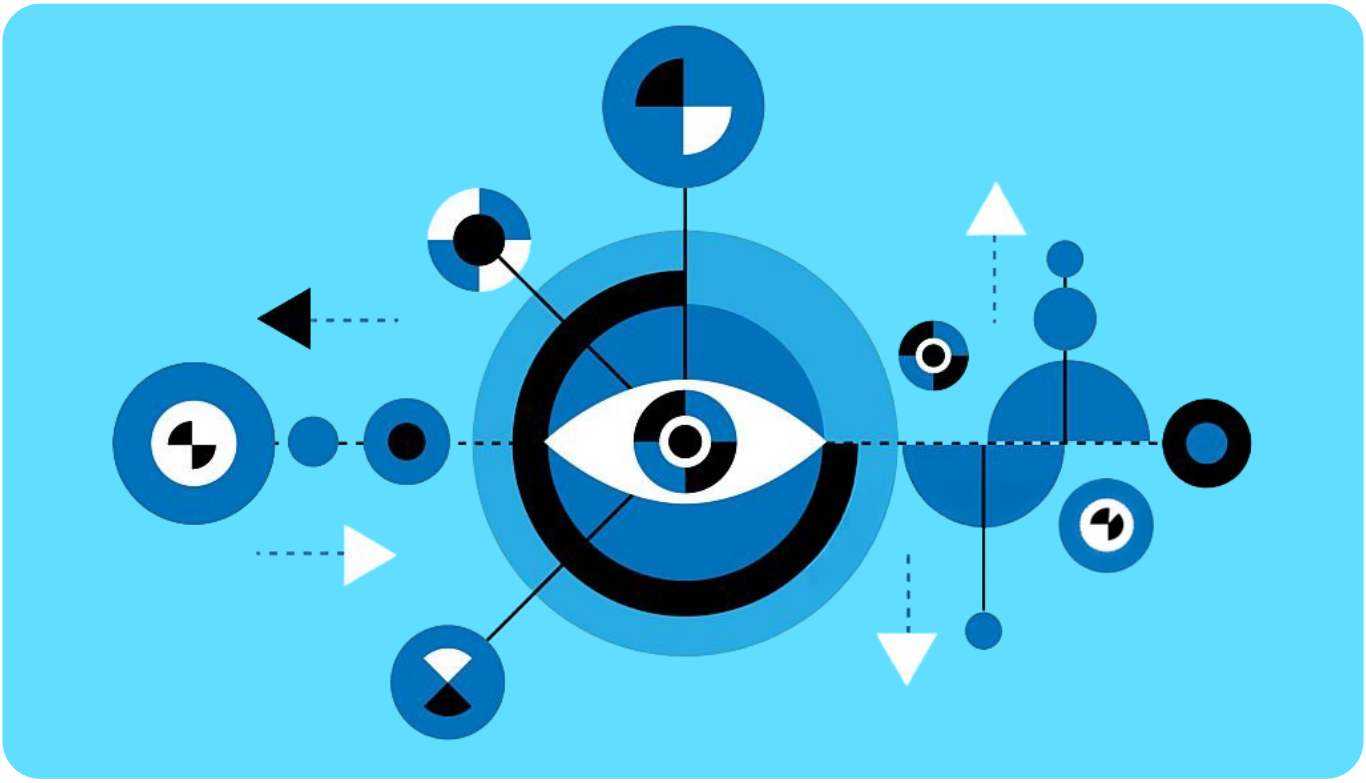

B. Transparency obligations

As it is explained in the preamble to the AI Act (recital 27), in general terms, transparency "means that AI systems are developed and used in a way that allows appropriate traceability and explainability". The EU AI Act is the first legislative framework that provides a regulation of transparency:

- Regarding content protected by copyright and related rights and used by GPAI providers to train their models, Article 53(1)(d) of the EU AI Act provides an obligation to "draw up and make publicly available a sufficiently detailed summary about the content used for training of the general-purpose AI model, according to a template provided by the AI Office." As it is explained in recital 107, such summary should balance the right of GPAI developers to "protect trade secrets and confidential business information" with the rights of rightholders and the enforceability thereof, including having regard to the applicability of relevant exceptions and the rights reservation possibility afforded by Article 4(3) of the DSM Directive (see above at §I.A.1).
- On 24 July 2025, the European Commission published a template for the public summary of training content for GPAI models. The template complements the guidelines on the scope of the rules for general-purpose AI models, published on 18 July 2025, and the General-Purpose AI Code of Practice released on 10 July 2025.

A transparency obligation in the terms described above has been only legally required in the EU. Nevertheless, calls to ensure the transparency of the data used for training purposes have been made in several other jurisdictions, including inter alia the UK.

Examples of jurisdictions where transparency requirements are currently being discussed include:

- **UK:** Amidst the ongoing potential reform discussion (see above at §I.A.1) and specifically pending the adoption of the Data (Use and Access) Act 2025, the Publishers Association commissioned research showing that 92% UK Members of Parliament believe that AI developers should be required to be transparent with authors and publishers about the works used to train their models, and 86% agree that compensation ensures the creative efforts of rightsholders are valued and protected (Publishers Association, MPs call for AI transparency and fair compensation for copyright content (1 April 2025)). The same view was shared by several authors, artists, and rightholders in an open letter to the UK Prime Minister on 10 May 2025 (Creative Titans and Industry Leaders Urge the Prime Minister to Support UK Creativity and Economic Growth by Enforcing Copyright Law (10 May 2025)). Despite all of this, a transparency obligation was regrettably not included in the adopted version of the Act, the justification being that this legislation differs and is separate from an official response to the Copyright and AI consultation, for which an official Government response is still pending at the time of writing.

In autumn 2025, rightholders in inter alia the visual sector called for retrospective settlements for past unauthorised use, transparent disclosure of training datasets and fair licensing agreements to ensure creators are properly credited and compensated for their contributions to AI development (DACS, Our Joint Statement Calling for Transparency, Fairness and Respect for Creators' Rights in the Age of AI (16 October 2025)). Still in autumn 2025, the British Copyright Council highlighted how respecting authors' and rightholders' copyright and related rights is also needed from a fundamental / human rights perspective (British Copyright Council, BCC Responds to the Call for Evidence by the Joint Committee on Human Rights (17 September 2025)).

- **USA:** No transparency obligation comparable to that of the EU AI Act exists under U.S. law, though a number of bills are currently pending, including The Transparency and Responsibility for Artificial Intelligence Networks (TRAIN) Act, which would require AI models to disclose training data records. That said, labelling practices relating to outputs are emerging, and so are related obligations (at least at state level): see below at §II.B.

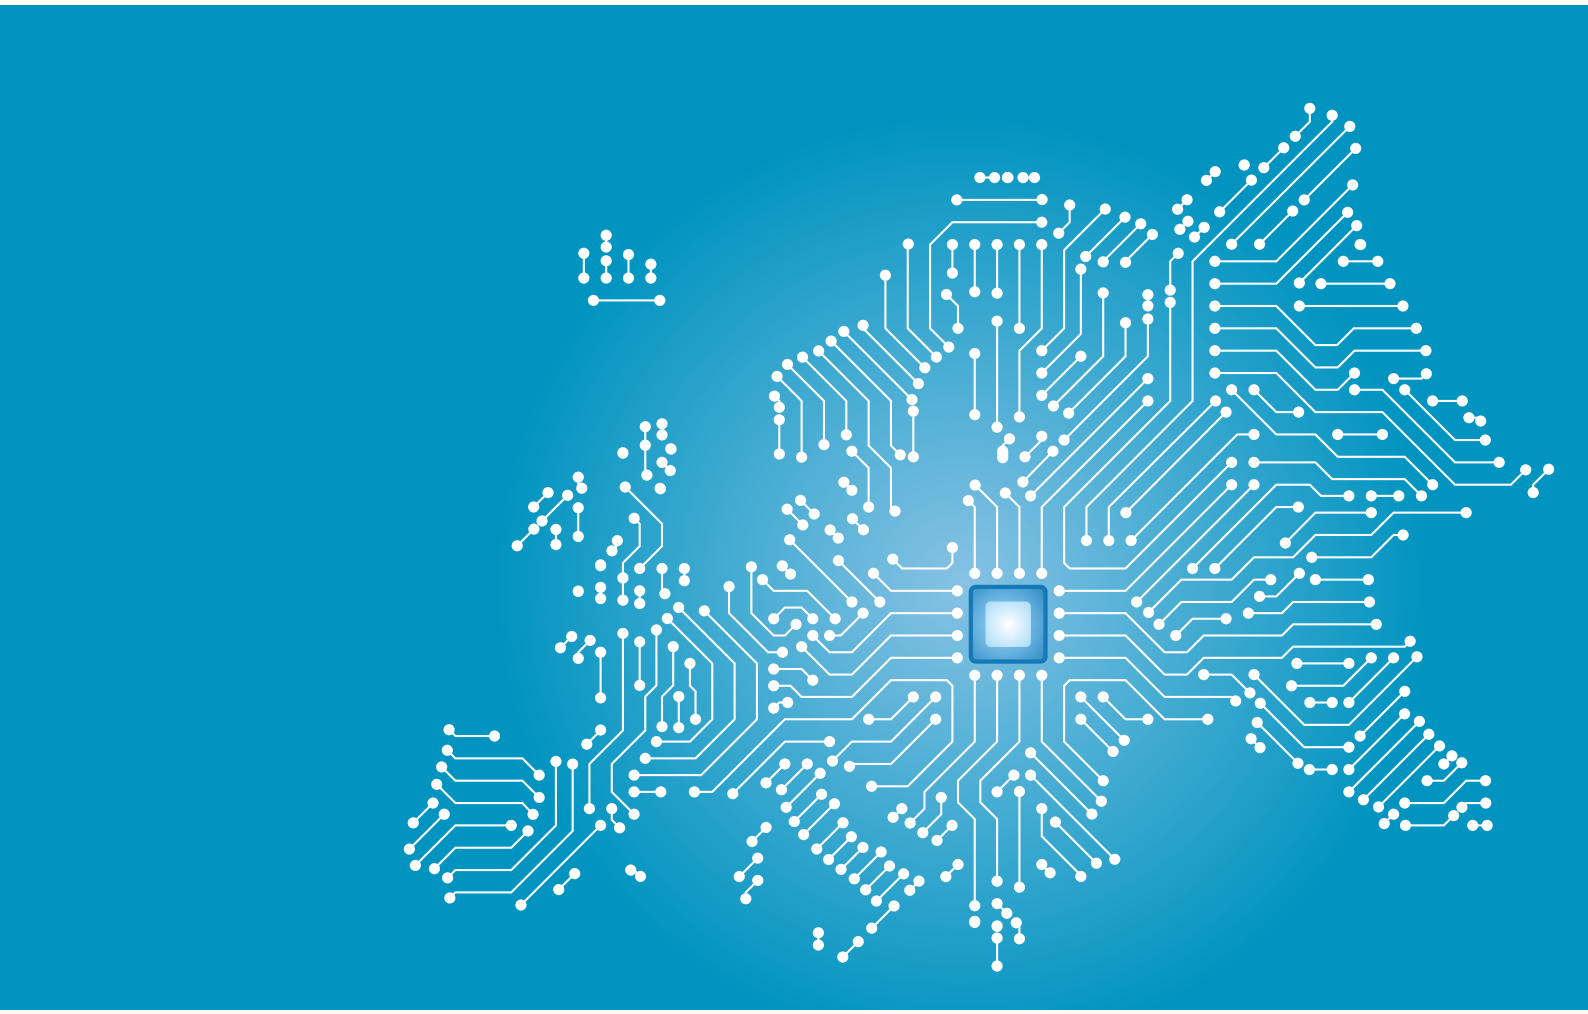

C. Synthetic media

The notions of synthetic media, digital replicas, and deepfakes refer to media content – including images, videos, audio, or text – generated or manipulated using technologies like but not necessarily limited to AI and intended to realistically but falsely depict an individual. The AI Act specifically defines a deepfake as an "AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful" (Article 3(60)).

From a legal perspective, the rise of synthetic media raises several issues, including but not limited to the **protectability of styles under copyright and, more broadly, forms of persona protection**. Above all, it raises the fundamental challenge of **how the creation and spread of non-consensual synthetic media may be prevented and repressed**.

Synthetic media is not new, but the rise of generative AI has exacerbated the creation and dissemination of non-consensual synthetic media, as well as replicas and imitations of one's own likeness and personal attributes, including voice and style. It takes very little time and effort to create a deepfake (S Bond, 'It takes a few dollars and 8 minutes to create a deepfake. And that's only the start' (23 March 2023) NPR).

Detecting synthetic media is key for several reasons, including to ensure reliability, provenance and integrity of scientific records and data. It also serves to contrast misinformation and manipulation (World Economic Forum, 'Why detecting dangerous AI is key to keeping trust alive in the deepfake era' (7 July 2025)). While challenges in this sense are widespread, tools are currently being developed and assessed to facilitate detection tasks (see, e.g., D Salvi and Others, 'Not all deepfakes are created equal: Triaging audio forgeries for robust deepfake singer identification' (20 October 2025)).

Once again, the policy and legislative discourse has been proceeding at an uneven pace:

- **EU**: Besides protection available under national laws (e.g., image/personality rights: see below at §I.C.2), data protection/privacy, as well as intellectual property rights, the AI Act requires deployers of an AI system that generates or manipulates image, audio or video content constituting a deepfake, to disclose that the content has been artificially generated or manipulated (Article 50(4)). It should be noted that no comparable legal obligation is foreseen for users of an AI system and/or other services, including but not limited to social media platforms. Nevertheless, some social media platforms have been implementing their own labelling policies: see below at §II:B.

  - **Italy**: Through the adoption of Law 132/2025, Italy has seemingly been the first EU Member State to introduce national legislation on AI. Among other things, the Law has criminalised the non-consensual dissemination of content generated by or altered through AI (Article 612-quater of the Italian Criminal Code).

- **China**: In 2025, the Cyberspace Administration of China issued the Labeling Measures for Content Generated by Artificial Intelligence, with effect as of 1 September 2025. The Measures relate to AI-generated synthetic content and target online information service providers (T Bond – E Ren, 'New AI content labelling rules in China: What are they and how do they compare to the EU AI Act?' (20 May 2025) Bird&Bird).

- **UK**: There is no general obligation under UK law comparable to that set in the AI Act. Specific provisions are nevertheless in place for the creation (Data (Use and Access) Act 2025, section 138) and dissemination (Online Safety Act 2023, sections 188 ff.) of non-consensual synthetic media of an intimate kind.

- **USA**: Besides protection available at the level of some states through publicity rights and the discussion surrounding the potential introduction of a federal publicity right (U.S. Copyright Office, Copyright and Artificial Intelligence – Part 1: Digital Replicas (July 2024)), the adoption of the No FAKES Act (reintroduced in 2025), which is meant to "protect intellectual property rights in the voice and visual likeness of individuals, and for other purposes" is pending. Meanwhile, the TAKE IT DOWN Act was adopted in 2025: it requires "covered platforms to remove nonconsensual intimate visual depictions, and for other purposes."

- **Brazil:** In autumn 2025, a bill that would amend Brazil's copyright law (PL 5005/2025) was introduced to protect the "image, voice, and other personal characteristics against realistic digital imitations generated by artificial intelligence or similar technology, as well as [to reinforce] the rights of performing artists".

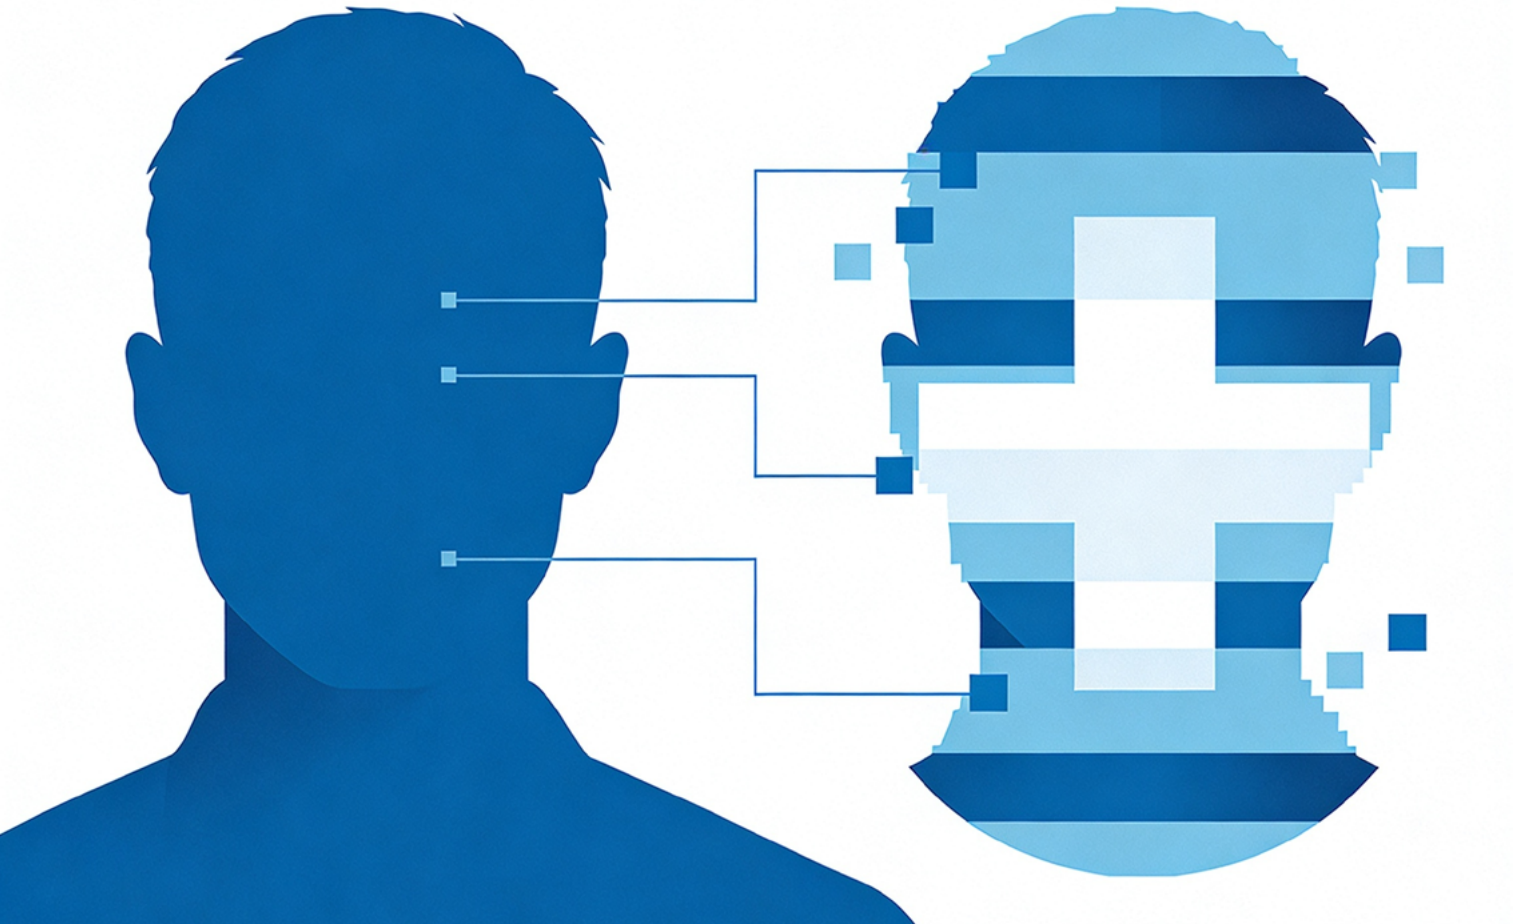

1. Protectability of styles under copyright

As noted, a key characteristic and goal of synthetic media is to imitate a person. From a copyright perspective, the possibility for an AI model to generate outputs that imitate someone's style has been not only widely reported by and discussed in general media – for example having regard to the possibility unveiled by OpenAI to inter alia generate content in the style of Studio Ghibli's creations through its 4o Image Generation service (see, e.g., A Reisner, 'ChatGPT turned into a Studio Ghibli machine. How is that legal?' (13 May 2025) The Atlantic) – but is also at the centre of currently pending lawsuits, including in the USA (Disney Enterprises, Inc., et al. v. Midjourney, Inc., C.D. Cal. Case No. 25-5275 (filed 11 June 2025); Warner Bros. Entertainment Inc, et al. v. Midjourney, Inc., C.D. Cal. Case No. 2:25-cv-08376 (filed 4 September 2025); Disney Enterprises, Inc., et al. v. Minimax, et al., C.D. Cal. Case No. 25-8768 (filed 16 September 2025)).

At the outset, it should be noted that the circumstance that an AI model may imitate someone's content indicates that the training data underpinning that model incorporates such content. This basic consideration is key not only to the analysis at §I.A above but also the transparency obligations discussed at §I.B above.

At output level, unauthorised style imitation raises the question whether it might be actionable under at least the right of reproduction. While the answer depends on the circumstances at hand, including (i) the applicability of the **idea/expression dichotomy** (Articles 2 of the WIPO Copyright Treaty and 9(2) of the TRIPS Agreement), (ii) proof of **derivation**, and (iii) whether the taking in question is of a **protected part, there is no reason to rule out in principle the actionability of style imitation, at least in certain cases.**

Insofar as copyright is concerned, to be protected, the style in question must be, first of all, a 'production' in a Berne sense, that is a 'work' that is also sufficiently 'original'.

EU: The CJEU has defined the notion of work as referring to an expression (as opposed to an idea) that is identifiable with sufficient precision and objectivity (CJEU, Levola Hengelo, C-310/17, para 40). For the notion of originality, see below at §II.A.

South Korea: The Guide to Preventing Copyright Disputes Related to Generative AI Outputs, released by the Ministry of Culture, Sports and Tourism and the Korea Copyright Commission in 2025, recalls that, under the Copyright Act, mere facts and ideas, not expressions, are not protected. Nevertheless, "[because the [border] between an idea and an expression may differ in each and every case, trying to distinguish between the two in general terms is difficult. * For example, a user directly references the work of a particular author or uses prompts to induce an output that is similar to a particular creative work, so that the GenAI output has expressions similar to an existing work. * For reference: An operator that uses works of a famous author without permission to train its AI model to mimic a particular style that it will then use in a business model or for publicity may be violating the Unfair Competition Prevention Act." p. 62).

As a further confirmation that a style might be protected under copyright and related rights there is the existence, e.g., under both EU and UK laws, of specific exceptions allowing **pastiche** (Article 5(3)(k) InfoSoc Directive; Article 17(7)(9) DSM Directive; section 30A CDPA).

EU: At the time of writing, the first CJEU referral specifically asking about such a notion is pending (CJEU, Pelham II, C-590/23). However, in his Opinion in Pelham I, C-476/17, Advocate General Szpunar suggested that pastiche "consists in the imitation of the style of a work or an author without necessarily taking any elements of that work" (fn 31; emphasis added). In his Opinion in Pelham II, C-590/23 Advocate General Emiliou admitted that, when someone "imitates closely the style of a single work", such an activity may be relevant under copyright's right of reproduction and the pastiche exception: "The elements borrowed, while 'stylistic', could still be regarded as original, especially when combined." (para 66)

UK: In Shazam Productions Ltd v Only Fools The Dining Experience Ltd & Ors (Rev1) [2022] EWHC 1379 (IPEC) (08 June 2022), Judge John Kimbell QC considered that the everyday meaning of pastiche inter alia entails the imitation of the style of pre-existing works and the use or assemblage of pre-existing works in new works (para 195).

While the idea/expression dichotomy remains a cornerstone of copyright and – to establish infringement – proof of derivation is required, in certain cases a style may be regarded a protected work. As such, **the reproduction in verbatim or altered form of such a style may be also relevant under the right of reproduction.**

2. Persona protection through different tools

Ensuring effective protection of one's personal attributes – including appearance, style, name, silhouette, voice (on the latter, see A Baris, 'AI covers: legal notes on audio mining and voice cloning' (2024) 19(7) Journal of Intellectual Property Law & Practice 571), etc – is a growing concern in the age of generative AI.

Litigation relating specifically voice misappropriation in the context of AI model development has been brought in a number of jurisdictions, including in the USA (Paul Lehrman, et al., v. Lovo, Inc., S.D.N.Y. 24-CV-3770 (JPO)), and India (Arijit Singh v. Codible Ventures LLP and Ors, on which see A Baris, 'Publicity rights in the AI era: Key takeaways from artist Arijit Singh's recent legal Victory in India' (31 August 2024) The IPKat).

> The legal tools available to protect personal attributes – including likeness and other personal features, e.g., voice, silhouette, name, style, etc. – **vary substantially** from country to country. In several instances, protection is possible under the umbrella of intellectual property rights. Some jurisdictions also specifically provide for standalone personality, image, or publicity rights.

Where personality, image, or publicity rights exist, there are often significant variations, including regarding:

1. Beneficiaries (only natural or also legal persons, if not – in some jurisdictions (e.g., Italy) – even cultural heritage artifacts (E Rosati, 'Bologna Court of Appeal confirms that Italian image rights extend to cultural heritage assets (which cannot be used commercially without authorization' (30 October 2024) The IPKat));

2. Scope (only actual and direct use of someone's likeness and other personal attributes, or also evocation);

3. Duration (life of the person, post-mortem protection, or even unlimited term);

4. Level of protection (in federal countries, country-wide or only state-level. The latter is currently the case in, e.g., the USA and has been tackled in the context of the rise of generative AI as a matter of policy in U.S. Copyright Office, Copyright and Artificial Intelligence – Part 1: Digital Replicas (July 2024));

5. Enhanced protection of certain persons (celebrities as opposed to lay people, children vis-á-vis adults), or settings and situations (private activities as opposed to public activities and events); and

6. Available remedies.

Copyright, performers' rights, and trade mark rights may be used to protect one's persona too. Traditional rights under copyright, including economic and moral rights may be used in some situations to prevent the unauthorised use of one's own image. In certain countries, copyright statutes further envisage protection of rights of portrayal or portrait rights.

Turning to trade mark law, decisions have been rendered in multiple jurisdictions regarding the registrability of signs and enforceability of resulting registrations that consist of one's face, signature, silhouette, name, etc (see recently B Fritz, 'Matthew McConaughey Trademarks Himself to Fight AI Misuse' (13 January 2026) The Wall Street Journal).

It is also important to note that specific instruments may be supplemented by other legal tools. For example, in jurisdictions where image/publicity rights are not explicitly recognised via standalone rights, protection of one's likeness and other personal attributes might be secured by recourse to other causes of action, e.g., passing off / unfair competition, breach of confidence / misuse of private information, libel / slander / defamation laws, as well as other civil or criminal law torts. Data protection and privacy laws may offer protection too, not only considering that several personal attributes qualify as personal data but also because, in some jurisdictions (including the EU), the law grants – at certain conditions – a 'right to be forgotten' (Everything You Need to Know about the "Right to Be Forgotten").

In all of this, **the international and regional level playing field remains minimal**. For example, across the various European countries, including those that are EU members, the diversity of available protections remains substantial. In the EU, the AI Act regulates **deepfakes** to a certain extent, also imposing labelling and disclosure obligations on the side of AI deployers (Article 50(4): see also below at §II.B). At the national level, proposals have been tabled in, e.g., Denmark (World Economic Forum, Deepfake Legislation: Denmark Moves to Protect Digital Identity (30 July 2025)) and the Netherlands (D Visser, 'The Dutch and Danish proposals for legislation on deepfakes' (18 July 2025) AI-forum) to introduce **new rights** that would allow persona protection against non-consensual synthetic media creation and dissemination. Outside of the regulatory and legislative context, there have also been **private efforts** to address synthetic media and personality rights, like for example TikTok's 'AI-generated' label (see also below at §II.B).

3. AI misappropriation in publishing

As noted above at §I.A, the quality of an AI model depends in turn on the quality of the training data, with studies noting that the use of human-generated content, as opposed to synthetic content, leads to higher-quality AI results. A recent study conducted by researchers at Stony Brook University and Columbia Law School (T Chakrabarty – JC Ginsburg – P Dhillon, 'Readers prefer outputs of AI trained on copyrighted books over expert human writers' (last revised 9 November 2025) Columbia Public Law Research Paper No. 5606570 has found that AI models fine-tuned on just two books can generate writing in the style of famous authors (see also J Kemper, 'AI models can mimic famous authors' writing styles using just two books for training' (26 October 2025) the decoder).

As discussed above, styles are protectable by inter alia copyright if they are original works. In turn, an author whose style is imitated by AI should have different enforcement options at their disposal. First, they should be able to know whether their works have been used in the training of an AI model – including through robust transparency obligations on the side of AI developers (see above at §I.B) – so that they can enforce relevant rights under copyright. Second, at output generation, they should be also in a position to enforce their rights, including those granted under copyright (through the regurgitation of training data: see, e.g.: E Rosati, 'Infringing AI: Liability for AI-generated outputs under international, EU, and UK copyright law' (2025) 16(2) European Journal of Risk Regulation 603; MP Goodyear, 'Artificial Infringement' (last revised 25 April 2025) (forthcoming) UC Law Journal) and the rights in their personal attributes, as discussed above at §I.C.2).

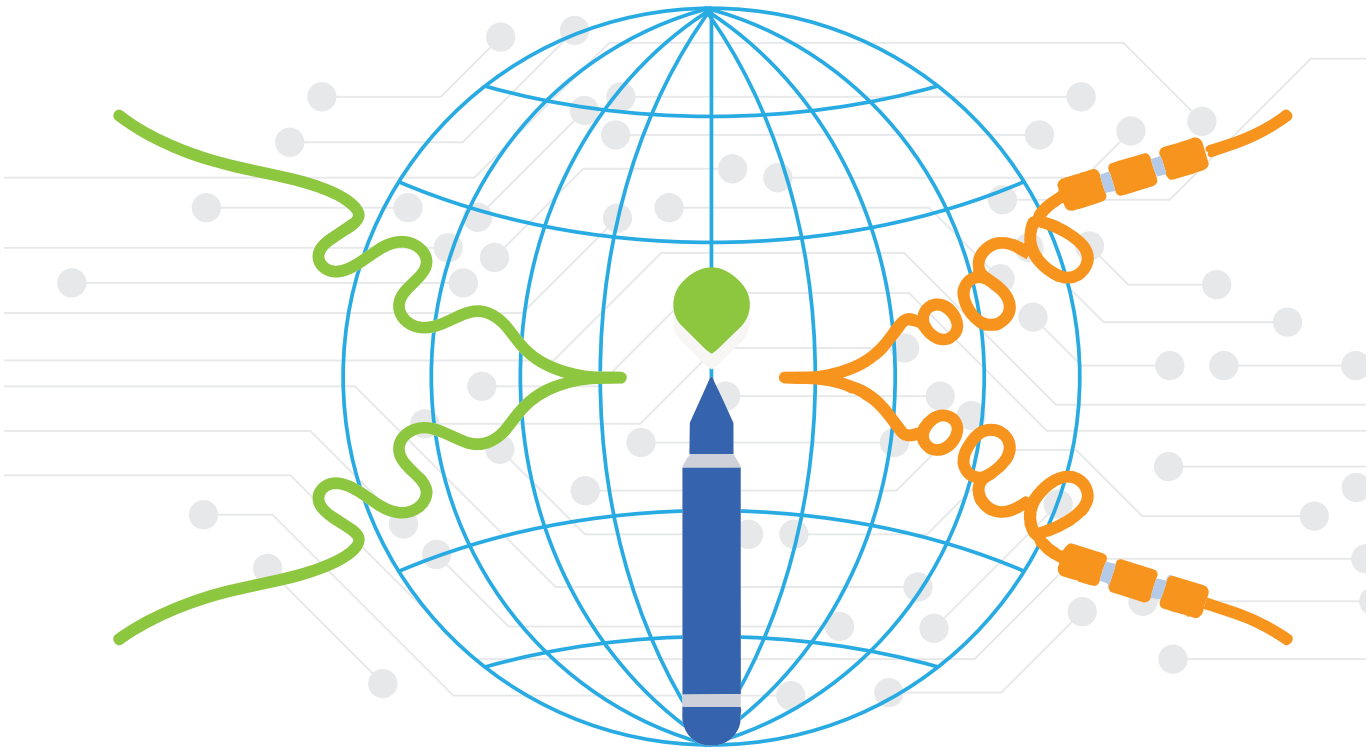

II. AI-GENERATED OUTPUTS

Turning to specific issues raised during the output generation phase, besides the potential liability connected to regurgitation of training data (also discussed in the IFRRO AI Toolkit) and the issues raised in connection with, e.g., the creation and dissemination of synthetic media (see above at §I.C), protectability under copyright law is a key aspect, and so are labelling requirements. Both are discussed below at §II.A and §II.B, respectively.

A. Human authorship requirement

The use and inclusion of AI-assisted and -generated components is growing across all industries, including creative ones, and so is the availability of AI-based tools and services that allow all of that.

The question whether the resulting content is protectable under copyright has been debated for a while. Most jurisdictions require that an author in a copyright sense is a human being. Examples include:

- **EU**: While the CJEU has not yet expressly ruled on the notion of author and EU directives do not provide a definition of this concept either, the Opinion of Advocate General Trstenjak Painer-145/10, para 121 is adamant that "only human creations are […] protected, which can also include those for which the person employs a technical aid, such as a camera").
  - **Germany**: In February 2026, a German first instance court confirmed that copyright is premised on human authorship (AG München, Endurteil vom 13.02.2026 - 142 C 9786/25).
- **UK**: Section 9(3) CDPA in the UK provides a legal fiction regarding computer-generated works, in the sense that the person who made the necessary arrangements shall be deemed to be the author. While some have questioned the compatibility of this approach with EU law, it should be noted that the approach under section 9(3) is likely to be regarded as a sui generis right, rather than copyright. Overall, it appears that a human author is required for copyright to arise, even under UK law.
- **USA**: The Compendium of the U.S. Copyright Office Practices, 3rd edn, §306, p. 7 is clear that only human-created works are registrable. In Thaler v. Perlmutter, No. 23-5233 (D.C. Cir. 2025) it was acknowledged that, while the U.S. Copyright Act does not define 'author', this notion only refers to human beings "[because many of the Copyright Act's provisions make sense only if an author is a human being", including inter alia ownership, duration, inheritance, and

transfer of copyright. Furthermore, the statute refers to machines as tools, not as authors. Above all, "adhering to the human-authorship requirement does not impede the protection of works made with artificial intelligence." That is so inter alia because a human can create works with the assistance of AI and those works can receive protection. In March 2026, the U.S. Supreme Court denied certiorari in Thaler v. Perlmutter, thus leaving the D.C. Circuit's ruling unaffected.

- **South Korea**: Article 2 of the Copyright Act requires a human author for copyright to arise, given that 'work' "means a creative work that expresses human thoughts and emotions" and an 'author' is "the person who has created a work". The Guide to Copyright Registration for Generative AI-Assisted Works, released by the Ministry of Culture, Sports and Tourism and the Korea Copyright Commission in 2025, provides criteria to determine when copyright is available to AI-assisted works (p. 12): "The human creative contribution can be acknowledged in the following cases: 1. when a GenAI output, generated by a user who inputs their copyrighted work as the prompt, reflects the creativity of that copyrighted work; 2. when the additional work performed by a user on an autonomous output—such as modification, addition, or deletion—displays creativity; 3. when creativity is present in the selection, arrangement, or composition of an autonomous output."

- **Singapore**: A position essentially identical to that adopted in South Korea has been endorsed by the Government of Singapore: An AI output may qualify for copyright protection "if there is sufficient engagement of the human intellect [and] the user exercises significant creative or editorial control."

On a practical level, the real question is thus likely going to be not whether non-humans may be authors, but rather to what extent one may use AI and still be regarded as the author of the resulting output. In other words, the issue is how much **human intervention** is required to sustain copyrightability of a human-authored work created with the assistance of an AI tool.

To answer this, it is necessary to construe the notion of authorship in light of another key requirement for copyright subsistence: originality. In several jurisdictions, originality requires some degree of creativity, though the actual level thereof may vary:

EU: The EU standard of originality entails a "'creative' aspect, and it is not sufficient that the creation of [the work] required labour and skill." (Opinion of Advocate General Mengozzi in Football Dataco, C-604/10, para 35; Opinion of Advocate General Szpunar in Funke Medien, C-469/17, para 18) The EU originality criterion is not satisfied when the creation of a work "is dictated by technical considerations, rules or constraints which leave no room for creative freedom" (CJEU, Cofemel, C-683/17, para 31; more recently, CJEU, Mio/konektra, C-580/23 and C-795/23, para 49) Overall, the EU originality requirement can be summarised as follows: "In order for an intellectual creation to be regarded as an author's own it must reflect the author's personality, which is the case if the author was able to express [their] creative abilities in the production of the work by making free and creative choices." (CJEU, Funke Medien, C-469/17, para 19).

UK: Even though historically the UK has envisaged a low standard of originality ('sufficient skill, labour or effort'), its tenure as an EU Member State has required UK courts to comply with the EU approach to originality as shaped by the CJEU (see above). Such an approach has been maintained post-Brexit (see e.g., THJ Systems Limited & Anor v Daniel Sheridan & Anor [2023] EWCA Civ 1354, para 15: "Section 1(1)(a) of the 1988 Act must, so far as possible, be interpreted in accordance with Article 2(a) of European Parliament and Council Directive 2001/29/EC of 22 May 2001 on the harmonisation of certain aspects of copyright and related rights in the information society ("the Information Society Directive") as interpreted prior to 31 December 2020 by the Court of Justice of the European Union.")

USA: A work of authorship must possess "some minimal degree of creativity" to sustain a copyright claim (Feist, 499 U.S. at 358, 362; see also Compendium of the U.S. Copyright Office Practices, 3rd edn, §308.2, p. 9). The U.S. Copyright Office's Copyright and Artificial Intelligence – Part 2: Copyrightability (January 2025) confirms that the use of AI models to assist human creativity is no bar for copyright protection of the output insofar as there is sufficient human control over the expressive elements. This conclusion is also consistent with the 1965 Report to the Librarian of Congress (p. 5).

South Korea: The law requires a human creative contribution. As it is explained in the Guide to Copyright Registration for Generative AI-Assisted Works, released by the Ministry of Culture, Sports and Tourism and the Korea Copyright Commission in 2025 (p. 13), "A creative expression refers to an expression that is not a copy of another's work, reflects human mental effort, and is distinguishable from other copyrighted works. Creativity does not mean originality in the complete sense; rather, it merely means that a work is not simply a copy of another's and embodies the author's own unique ideas or emotions. To satisfy this requirement, it is sufficient that the work bears the characteristics of the result of the author's own mental effort and is distinguishable from existing works of other authors. (Supreme Court, 27th February 2014, Decision 2021da28745)".

Against the background illustrated above, **the author is the person who controlled the process resulting in the creation of the work in question by making the choices conferring it sufficient originality.**

All of this has been also applied in the few cases decided so far regarding AI-generated and -assisted outputs:

USA: The U.S. Copyright Office has allowed registration of works created in part with the aid of AI: see, e.g., Zarya of the Dawn (registration #VAu001480196) and A Single Piece of American Cheese (registration #VAu001543942). In Thaler v. Perlmutter, No. 23-5233 (D.C. Cir. 2025), the U.S. Court of Appeal for the District of Columbia Circuit reinstated that the notion of 'author' only refers to human beings. Overall, a human can create works with the assistance of AI, and those works can receive protection. The court remained unpersuaded that a requirement of human authorship would disincentivise creativity by creators and operators of AI.

China: Case law indicates that, if it is possible to identify a sufficient human contribution, AI-assisted outputs are protectable under copyright: see, e.g., Beijing Internet Court in Li v Liu, Civil Judgment (2023) Jing 0491 Min Chu No. 11279 and People's Court of Zhangjiagang City, Jiangsu Province in Butterfly Art Chair (Judgment of 19 March 2025 – (2024) Su 0582 Min Chu No. 9015; see T He, 'Distinguishing copyrightable from non-copyrightable AI-generated content (2025) GRUR Int).

The decisions referred to above do not indicate that AI-generated outputs are protectable tout court: instead, they stress the need to reconstruct the process that led to the AI-generated output in order to **identify the human author's contribution** and separate that from what is purely AI.

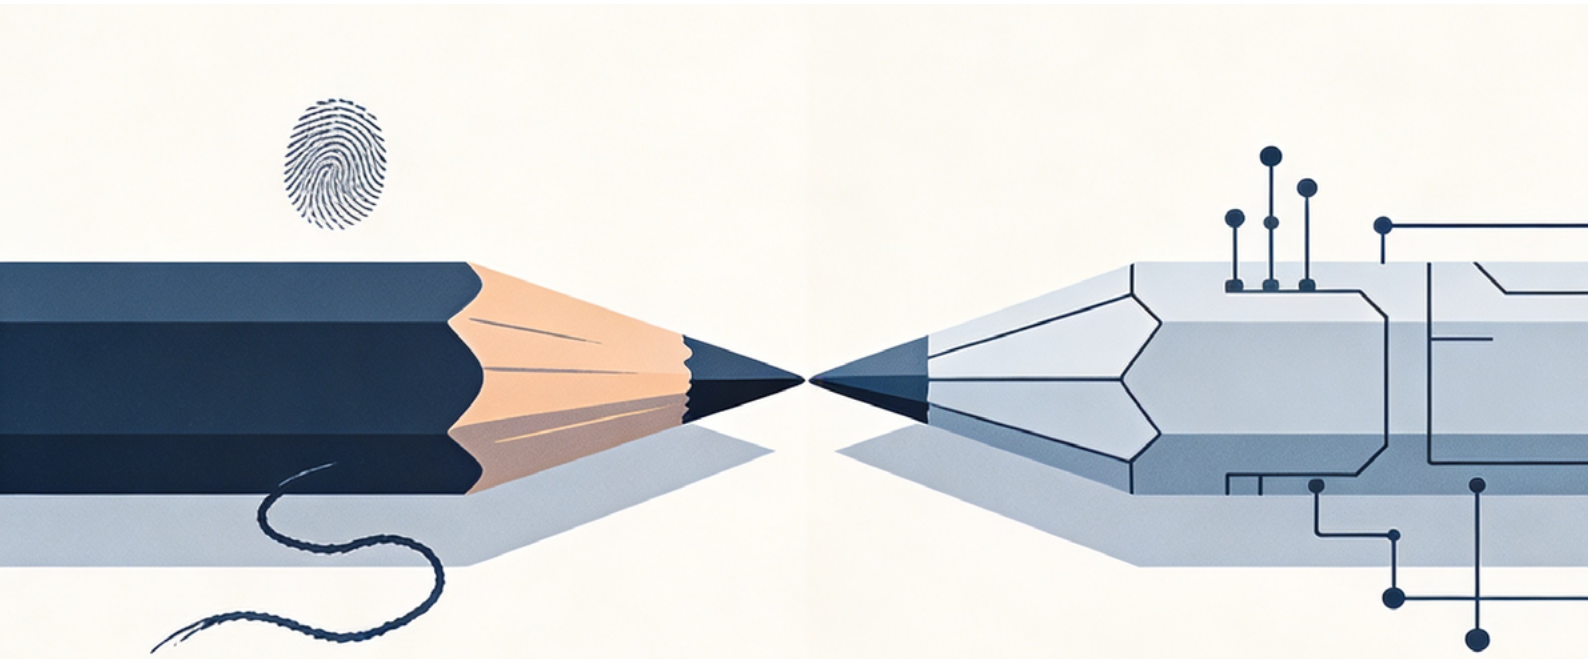

B. Labelling requirements

Besides transparency obligations regarding training data (see above at §I.B), some jurisdictions also provide for transparency obligations regarding the origin of content at output level, specifically requiring disclosure of its AI-generated nature by labelling it as such. Such obligations are foreseen in, e.g.:

- **EU**: Article 50(2) of the EU AI Act requires providers of AI systems, including GPAI ones, that generate synthetic content to label the resulting outputs as such in a machine-readable format. Such an obligation does not apply where AI systems are inter alia used with an assistive function for standard editing or do not substantially alter the input data provided by the deployer or the semantics thereof. Article 50(4) provides for specific obligations relating to deepfakes (see further recital 134 in the preamble to the AI Act; see also above at §I.C) and content of an informative kind. Article 50(7) further provides that the AI Office shall encourage and facilitate the drawing up of codes of practice at the EU level to enhance the effective implementation of these obligations and entitles the European Commission to adopt implementing acts to approve those codes of practice (see also recital 135 in the preamble to the AI Act).

- **USA**: While legislation has been suggested but not adopted at the federal level, the California AI Transparency Act of 2024 (California SB 942) mandates upon AI developers to embed imperceptible indications in AI-generated content, also entitling users to have such information displayed.

- **China**: As noted above at §I.C, the Cyberspace Administration of China has issued the Labeling Measures for Content Generated by Artificial Intelligence, with effect as of 1 September 2025. The Measures relate to AI-generated synthetic content and target online information service providers.

- **South Korea**: The Basic Act on the Development of Artificial Intelligence and the Establishment of Trust (adopted in 2025 and with effect from 22 January 2026) inter alia requires AI business operators providing products or services using high-impact AI or generative AI to notify users in advance that the product or service is AI-based and/or -generated. Such an obligation further extends to synthetic media.

Other jurisdictions are also considering legislating in this sphere, an example being **Brazil**, where Bill No. 2338/2023 proposes introducing labelling obligations in its current Article 19. At the time of writing, no country in the LATAM region appears however to have legislation in place imposing such an obligation.

In addition to legislative interventions, labelling technologies and standards are also being developed and deployed on a voluntary basis. Amongst these, the Coalition for Content Provenance and Authenticity (C2PA) provides an open technical standard for publishers, creators and consumers to establish the origin and edits of digital content. The current membership is broad and diverse. Furthermore, some companies are developing their own technologies: examples include Google's SynthID (currently in testing phase) and TikTok's 'AI-generated' label.

III. GOVERNANCE MODELS AND RECOMMENDATIONS

The discussion in the preceding Parts has highlighted how, on the one hand, the rise of generative AI has been raising similar – when not identical – issues across several jurisdictions. On the other hand, it has also shown that the discussion and related solutions have not reached the same maturity everywhere. While some jurisdictions have adopted specific legislation, others (the majority) have not yet intervened and it is currently unknown whether any legislative intervention will be ultimately undertaken and, if so, what form it will take.

Based on the analysis undertaken, the main recommendations for an AI governance model are as follows:

A. From AI training to output generation

- The training of AI models requires access to and use of large quantities of works and other subject-matter. Oftentimes such content is included in works protected by copyright and/or subject-matter protected by related rights. This is so for several reasons, including quality and timeliness of such content.

- The exclusive rights granted by copyright are proprietary and preventive in nature. That means that, to undertake any restricted act, prior authorisation is required unless an exception is available.

- In any event, there is no exception to copyright and/or related rights that fully covers unlicensed AI training in each and every instance, nor could there be. This is also due to the requirements of lawful access and the three-step test (UK: fair dealing): where, e.g., TDM exceptions exists, they only encompass specified acts undertaken for TDM purposes. AI training engages the doing of several restricted acts, with the result that in most cases authors' and rightholders' authorisation to train an AI model on their protected content is required. Such a conclusion should also apply having regard to fair use (open-ended) doctrines, especially but not necessarily solely where there is a licensing market for the training materials and where the resulting model is exploited commercially and generates outputs that compete with / supplement the training data, thus competing with them on the same market.

Licensing, not exceptions, is the sustainable and fair manner for AI to develop. Licensing solutions are currently emerging on the market, having regard to both direct licences and collective licences.

- AI developers should be transparent about the content used for training purposes, subject to using it having obtained prior authorisation or being eligible for the application of an exception (if available).

Legislatures should impose detailed transparency obligations on AI developers, and make it easy for rightholders effectively to protect their rights and prevent / request compensation for the unauthorised use of their content.

- The production and dissemination of synthetic media has been enhanced by the rise of generative AI. Currently, the international level playing field is thin and the regional (i.e., EU) one is not particularly developed. Labelling requirements exist in some jurisdictions though.

- The unauthorised creation and dissemination of synthetic media meant to resemble an actual person and/or their personal attributes – including their voice and style – could be contrasted through different rights and tools, including but not limited to personality / image rights, intellectual property rights, data protection / privacy laws, and other causes of action. Some countries are currently considering introducing new rights to contrast such a phenomenon.

- Insofar as styles specifically are concerned, they may also receive protection under copyright laws, provided that the style in question qualifies as an original work.

- Authors and rightholders should be able to enforce their rights against the use of their works during training and through imitation at output generation.

While persona protection may be achieved in different ways and through different tools, what is key is the ability to enforce such available protections effectively, so to contrast the creation and dissemination of content that is meant to realistically but falsely depict a living or deceased person and/or their personal attributes, including voice and style.

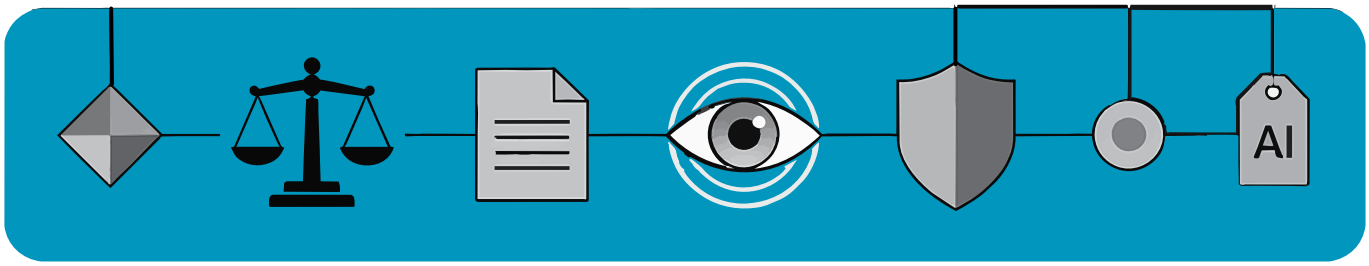

B. AI-generated outputs

- Most laws around the world require, for copyright to arise, that a work is created by a human author. In the context of AI-generated and -assisted outputs, this means that AI in itself cannot be regarded as an author.
- What is key is thus to identify the human contribution in AI-generated and -assisted works, so to determine if such a contribution may be considered authorial, including having regard to the originality thereof.

While human-created work needs to be safeguarded, it might be possible for AI-generated and -assisted outputs to receive protection, but such protection shall be limited to the parts that have been authored by a human person, who has created the work with the aid and support of AI.

- Some legislatures have recently introduced obligations of transparency (labelling) regarding the origin of content generated through AI. This inter alia serves user and consumer trust and broader democracy goals. It also supports the value of human-, as opposed to AI-, generated content. It should be noted that in some cases, labelling obligations do not extend to users of AI models.
- In parallel with and irrespective of legislative intervention, labelling practices and standards are being also implemented by private companies and bodies/organisations.

It is important to safeguard the trust of users and consumers, as well as to uphold the value of human-generated content. As such, effective labelling – including by means of generally applicable legal obligations – should be imposed in relation to AI-generated content. Companies' own labelling practices are also helpful and should be incentivised.

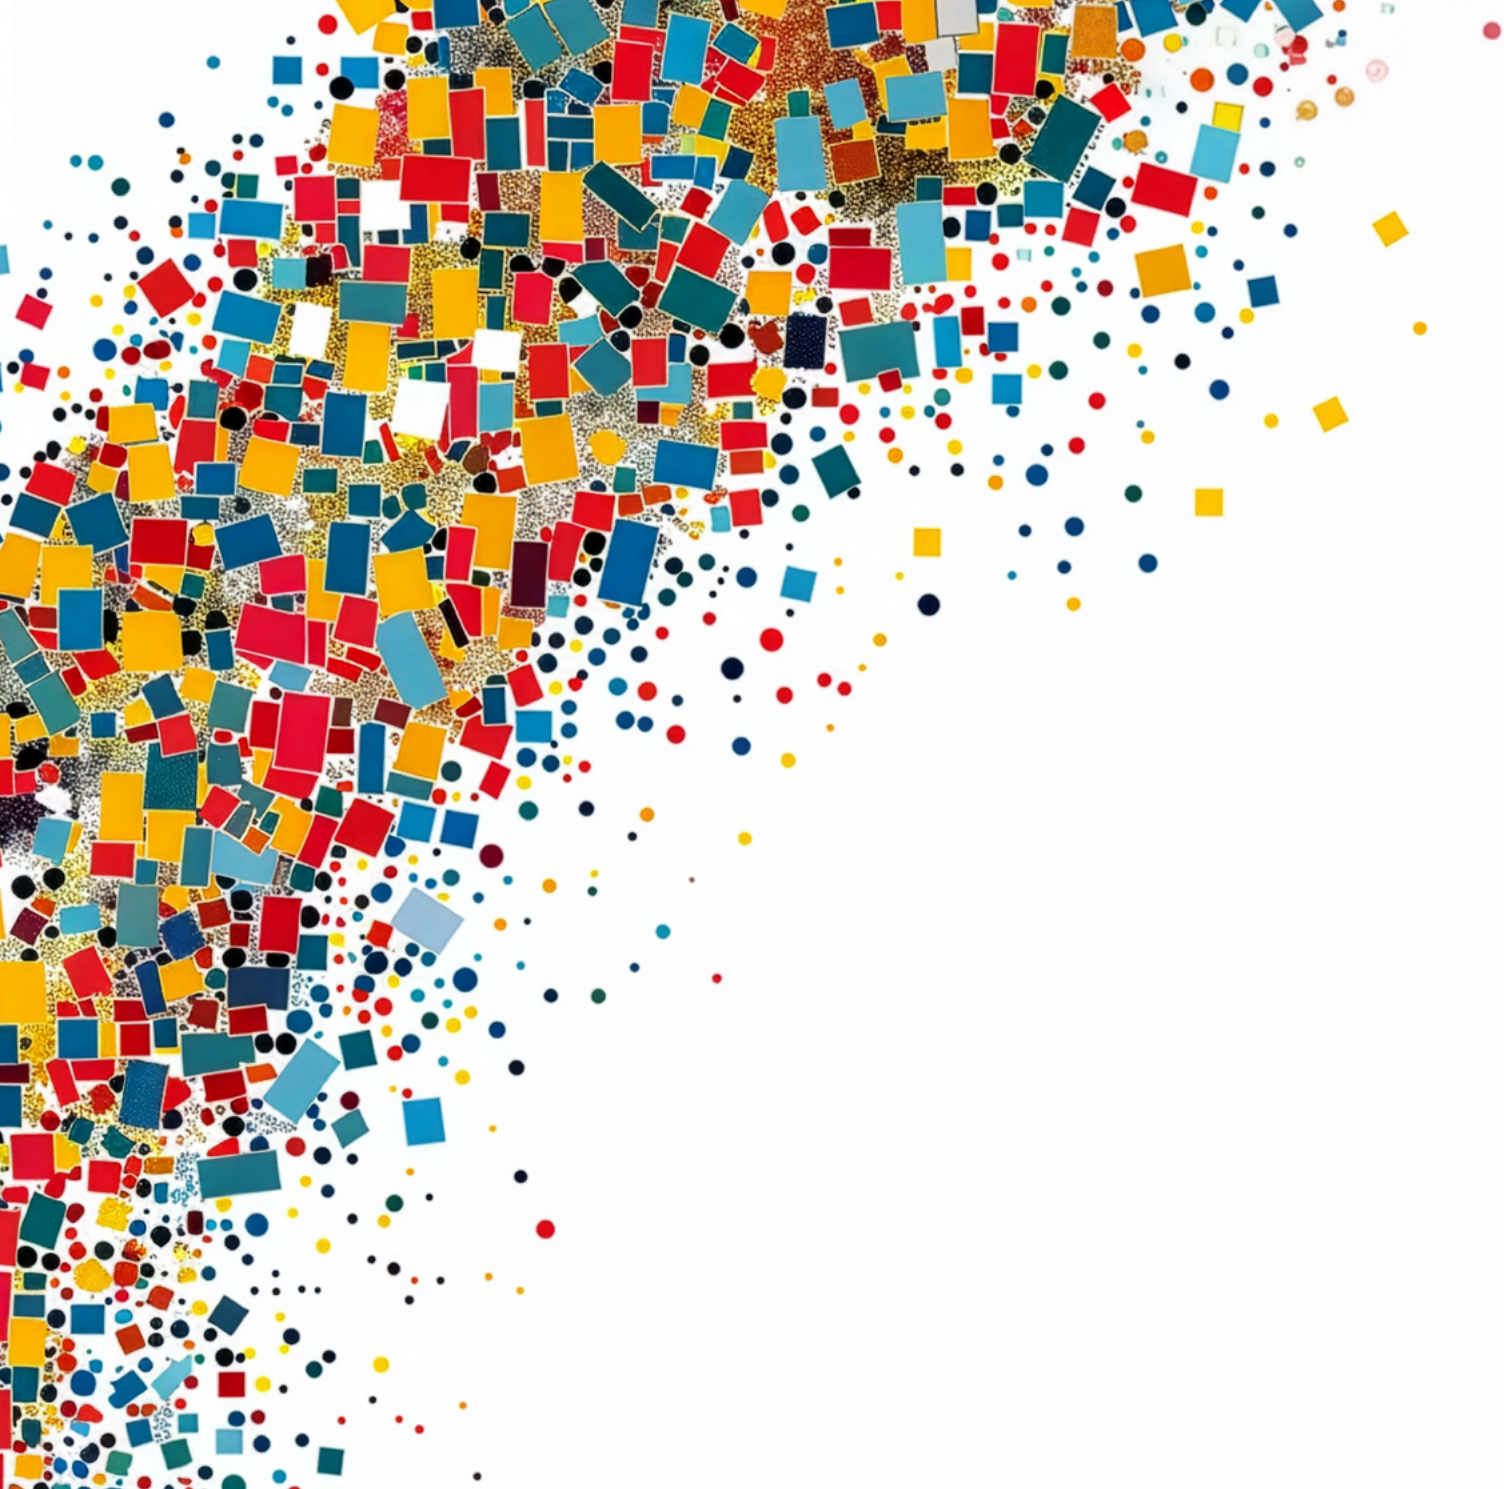

International Federation of Reproduction Rights Organisations (IFRRO)
Rue du Prince Royal 85-87 - 1050 Brussels, Belgium
T +32 2 234 62 60
secretariat@ifrro.org

Version 1.0 - May 2026

ifrro International Federation of Reproduction Rights Organisations