BILSEMMUN'26

WIPO
Study Guide

Co-Under Secretaries General

Beyzanur Özsığınan & Yağmur Raife Apaydın

Agenda Item:

Reconciling Generative Artificial Intelligence with the Global Intellectual Property Framework

Table of Contents

1. Letter from the Secretary-General

Dear Delegates,

It is with great pride and genuine excitement that I welcome you to BILSEM-MUN 2026. As Secretary-General, I have the honour of opening the doors of this conference to each of you whether you are stepping into a committee room for the very first time or returning as a seasoned diplomat.

Model United Nations is far more than a simulation. It is a space where ideas are tested, where the courage to speak meets the discipline to listen, and where young people discover that the world's hardest problems are rarely solved by a single voice. Over these days you will negotiate, you will disagree, you will compromise, and I hope you will surprise yourselves with what you are capable of.

This study guide is your starting point, not your finish line. Read it closely, question it, and then go further: study your country's position, understand the perspectives you will meet across the table, and arrive ready not merely to defend an argument but to build something with the people around you. The most memorable moments in any committee come from delegates who prepared deeply and then dared to think on their feet.

To our committee directors and the entire BILSEM-MUN team, thank you for the countless hours that make a conference like this possible. And to you, our delegates: this conference is yours. Fill it with sharp debate, bold ideas, and the kind of diplomacy that reminds us why we gather here in the first place.

I look forward to meeting you all and to the gavel that begins it all.

Warm regards,

Emir Mücteba Özsığınan
Secretary-General, BILSEM-MUN 2026

2. Letter from the Under-Secretary-General

Distinguished delegates of BILSEMMUN'26,

It is our greatest pleasure to welcome you to this committee. We are Beyzanur Özsığınan and Yağmur Raife Apaydın, your Co-Under-Secretary-Generals.

None of this would have been possible without the incredible work of the entire BILSEMMUN'26 team, who have poured their hearts and souls into every detail of this conference. We would especially like to extend our deepest gratitude to our Secretary-General, Emir Mücteba Özsığınan, whose vision, leadership, and dedication have made BILSEMMUN'26 the magnificent event it is today.

We cannot wait to witness what you will build together in committee.

Sincerely,

Beyzanur Özsığınan - Yağmur Raife Apayadın

Co-Under-Secretary-Generals

---

3. Introduction to the Committee

The World Intellectual Property Organization is the specialised agency of the United Nations responsible for intellectual property (IP). It was established by the WIPO Convention, signed in Stockholm in 1967 and entered into force in 1970, though its institutional roots reach back to the nineteenth century and the secretariats created to administer the Paris and Berne Conventions. Headquartered in Geneva, WIPO has 193 member states and administers more than two dozen international treaties spanning patents, trademarks, industrial designs, copyright and related rights, geographical indications, and — most recently — genetic resources and associated traditional knowledge.

WIPO's mandate is to promote the protection of intellectual property throughout the world and to encourage creative and inventive activity. It performs this mandate through several functions. It serves as the principal global forum for the negotiation of new IP norms and treaties. It administers global registration and filing systems — most importantly the Patent Cooperation Treaty (PCT) for patents, the Madrid System for trademarks, and the Hague System for industrial designs — that allow applicants to seek protection across many jurisdictions through a single

procedure. It provides capacity-building and technical assistance to developing countries. And it operates the WIPO Arbitration and Mediation Center for the resolution of IP disputes.

A point delegates must internalise from the outset: WIPO is a norm-setting and convening body, not a legislature or a court. It cannot compel a member state to change its copyright statute, and it does not decide the cases — Thaler v. Perlmutter, Bartz v. Anthropic, and the others described in this guide — that are shaping the law in real time. Those national rulings are evidence of the problem WIPO must address; they are not WIPO instruments. The realistic outputs available to this committee are therefore soft-law principles, model provisions, study and observatory mandates, harmonisation frameworks, and, at the most ambitious, the launch of negotiations toward a binding instrument. Resolutions that purport to bind national courts or override sovereign legislation will read as legally illiterate; resolutions that work with the grain of WIPO's actual powers will not.

WIPO has already built considerable machinery on this specific agenda. Since 2019 it has convened the WIPO Conversation on Intellectual Property and Frontier Technologies, the leading global multistakeholder forum on AI and IP, which over its first five years drew roughly 17,000 participants from 172 countries (WIPO 2025a). Its twelfth session, held in October 2025, focused on synthetic media, and in March 2026 WIPO launched the AI Infrastructure Interchange (AIII), a new initiative for global dialogue on copyright and AI (WIPO 2026). WIPO has also published practical guidance, including Generative AI: Navigating Intellectual Property and an IP policy toolkit for preparing innovation ecosystems for AI (WIPO 2024a). This committee builds directly on that body of work.

---

4. Introduction to the Agenda

Generative artificial intelligence — systems that produce text, images, audio, video, code, and even candidate molecular structures in response to prompts — has advanced from research curiosity to mass-market technology with extraordinary speed. The release of large language models and image generators to the public from late 2022 onward created hundreds of millions of users within months and, with them, a collision between a nineteenth- and twentieth-century legal framework and a technology that framework never anticipated.

The intellectual property system rests on a bargain that is centuries old. Society grants creators and inventors a set of exclusive, time-limited rights so that they have an incentive to create and disclose, and in exchange the public gains eventual access to the work and to the knowledge it contains. Copyright protects original expression; patents protect novel, non-obvious, useful

inventions; both assume, in ways that run deep through their statutes and case law, a human author or inventor at the origin. Generative AI strains both halves of this bargain simultaneously, and it does so from two directions at once.

At the output end, AI systems now generate material that looks, for all practical purposes, like the kind of expressive and inventive work the system was built to protect — but with the human role displaced, diffused, or reduced to a prompt. This raises the question of whether IP rights attach to AI-generated material at all, and if so, to whom. At the input end, these systems learn by ingesting vast quantities of existing protected works — books, images, music, code, news — overwhelmingly without licences. This raises the question of whether such training infringes the rights of the creators whose work fuels the machine, and whether existing exceptions such as fair use or text-and-data-mining can or should cover it. A third, fast-growing cluster of problems concerns synthetic media: deepfakes, voice clones, and digital replicas that implicate not only copyright but personality, publicity, and human-rights interests.

Cutting across all of this is a distributional question that WIPO, more than any other forum, must take seriously. The compute, capital, and data required to build frontier models are concentrated in a handful of firms and a handful of countries. Rules written for that frontier risk entrenching it, leaving low- and middle-income economies as consumers rather than producers of AI, and exposing the cultural output and traditional knowledge of communities in the Global South to extraction without consent or benefit. A framework that protects only rightsholders, or only developers, or only the wealthiest jurisdictions, will not hold. The task before this committee is to find language that keeps the incentive-and-access bargain functioning across a fractured and fast-moving landscape — without breaking it.

---

5. Key Terms and Definitions

Intellectual Property (IP): Legal rights arising from creations of the mind, including inventions, literary and artistic works, designs, and symbols, names and images used in commerce.

Copyright: The exclusive rights granted to the author of an original literary, artistic, musical, or other creative work to reproduce, distribute, adapt, perform, and display it, generally for the life of the author plus a set term.

Patent: An exclusive right granted for an invention that is new, involves an inventive step (is non-obvious), and is capable of industrial application, in exchange for public disclosure of the invention.

Author: In copyright, the creator of an original work; the status from which ownership and the bundle of exclusive rights flow.

Inventor: In patent law, the natural person or persons who conceived an invention. Distinct from the owner or assignee of the patent.

Originality: The requirement that a work originate with its author and possess at least a minimal degree of creativity in order to qualify for copyright protection.

Human-Authorship Requirement: The principle, applied in many jurisdictions, that copyright protection is available only for works created by a human being.

Generative AI: AI systems that produce new content (text, images, audio, video, code, structures) by learning statistical patterns from large training datasets and generating outputs in response to prompts.

Foundation Model / General-Purpose AI (GPAI) Model: A large model trained on broad data and adaptable to many downstream tasks; the EU AI Act's "GPAI model" is the regulatory unit to which transparency obligations attach.

Training Data: The corpus of works and other material on which an AI model is trained; the subject of the "input problem."

Fair Use: A flexible defence in U.S. copyright law permitting certain unlicensed uses, judged on four factors: the purpose and character of the use (including whether it is transformative and commercial); the nature of the work; the amount used; and the effect on the market for the original.

Fair Dealing: A more circumscribed, category-based equivalent of fair use found in the UK and many Commonwealth jurisdictions.

Transformative Use: A use that adds new expression, meaning, or purpose to the original; a central consideration under the first fair-use factor.

Text and Data Mining (TDM): The automated computational analysis of large bodies of text and data to identify patterns; in the EU, the legal category most directly covering AI training.

Opt-Out / Rights Reservation: A mechanism, ideally machine-readable, by which a rightsholder signals that their works may not be used for TDM or AI training.

Transparency Obligation: A legal requirement that an AI developer disclose information about the data used to train a model.

Compulsory / Statutory Licence: A mechanism permitting use of a work without the owner's consent in exchange for a set, regulated fee; frequently proposed as a middle path for AI training.

Collective Licensing: Licensing administered by a collecting society on behalf of many rightsholders, reducing transaction costs.

Public Domain: The body of works not protected by IP rights, freely usable by anyone; where purely AI-generated output falls in jurisdictions denying it protection.

Derivative Work: A work based upon one or more pre-existing works; relevant to whether AI outputs infringe the works they were trained on or prompted to imitate.

Digital Replica / Synthetic Media: AI-generated audio, image, or video; a digital replica reproduces a real person's voice or likeness.

Deepfake: Synthetic media that realistically depicts a real person doing or saying something they did not.

Right of Publicity / Personality Rights: The right of an individual to control the commercial use of their name, image, voice, and likeness.

Moral Rights: An author's non-economic rights, including the right of attribution and the right to the integrity of the work; protected under the Berne Convention.

National Treatment: The Paris and Berne principle that a country must grant foreign rightsholders the same protection it grants its own nationals.

Berne Convention: The foundational 1886 treaty on copyright, establishing automatic protection without formalities and minimum standards of protection.

Paris Convention: The foundational 1883 treaty on industrial property (patents, trademarks, designs), establishing national treatment and the right of priority.

Patent Cooperation Treaty (PCT): The WIPO treaty providing a unified international patent-filing procedure.

TRIPS Agreement: The WTO's 1994 Agreement on Trade-Related Aspects of Intellectual Property Rights, setting minimum IP standards binding on WTO members.

WIPO Copyright Treaty (WCT): The 1996 treaty extending copyright protection to the digital environment.

Traditional Knowledge (TK) / Traditional Cultural Expressions (TCEs): Knowledge, know-how, and cultural expressions developed and maintained by indigenous peoples and local communities.

Sui Generis Right: A right of its own special kind, created to address a problem existing categories do not fit; sometimes proposed for AI outputs or digital replicas.

Provenance / Watermarking: Technical measures that label content as AI-generated or trace its origin.

Capacity Building: Training and technical assistance to strengthen the IP institutions and expertise of developing-country members.

---

6. Historical Background

6.1 The bargain at the heart of intellectual property and the foundational treaties

Modern intellectual property law grew out of a simple economic intuition: creative and inventive work is costly to produce but cheap to copy, so without some legal mechanism to let creators capture a return, society will under-produce it. The answer most legal systems converged on was a temporary, exclusive right — a copyright or a patent — granted in exchange for creation and, in the case of patents, public disclosure. The right is deliberately limited in time precisely because the public interest in access is the other half of the bargain.

By the late nineteenth century the problem had become international: a work or invention protected in one country could be freely copied across the border. Two foundational multilateral treaties, both eventually brought under WIPO's administration, addressed this. The Paris Convention for the Protection of Industrial Property (1883) covered patents, trademarks, and

industrial designs, and established two enduring principles: national treatment, under which each member state must treat foreign applicants as it treats its own nationals, and the right of priority, which gives an applicant a window to file in other member states while preserving the original filing date. The Berne Convention for the Protection of Literary and Artistic Works (1886) did the same for copyright, and added principles that bear directly on this agenda: protection arises automatically upon creation without any registration formality, and member states must provide certain minimum standards, including a minimum term and protection for the author's moral rights — the rights of attribution and integrity.

These instruments were built around a human creator. Berne speaks of "authors" and protects works that are the expression of their intellectual creation; the entire architecture assumes a person at the origin. That assumption was uncontroversial for over a century. Generative AI is the first technology to make it genuinely contestable.

6.2 The human-authorship tradition: from Sarony to the monkey selfie

The principle that copyright protects the original expression of a human author has been reinforced repeatedly by national courts, well before AI arrived. In the United States, Burrow-Giles Lithographic Co. v. Sarony (1884) confronted the then-novel question of whether a photograph — produced by a machine, the camera — could be copyrighted. The Supreme Court held that it could, precisely because the photograph reflected the human author's creative choices in posing, lighting, and composition; authorship inhered in the human mind directing the machine, not in the machine. A century later, Feist Publications v. Rural Telephone Service (1991) confirmed that originality requires a minimal degree of creativity and rejected the "sweat of the brow" doctrine, holding that mere labour in compiling facts does not earn copyright. And in Naruto v. Slater (9th Cir. 2018) — the celebrated "monkey selfie" case — the courts confirmed that a non-human (there, a macaque that had tripped a camera shutter) cannot hold a copyright.

This lineage matters because it frames the AI debate not as a sudden rupture but as the latest test of a long-standing principle: the IP system has always located authorship and inventorship in the human, and has always treated tools — however sophisticated — as instruments of human creativity rather than creators in their own right. The central question generative AI poses is whether a system that generates output with little or no human intervention can still be fitted into that frame, or whether it falls outside it entirely.

6.3 The digital turn: TRIPS and the WIPO Internet Treaties

The late twentieth century brought the first major stress test of the framework: digitisation and the internet. Two developments are relevant background. First, the Agreement on Trade-Related

Aspects of Intellectual Property Rights (TRIPS), concluded in 1994 as part of the agreements establishing the World Trade Organization, set binding minimum standards of IP protection for all WTO members and linked IP to the trade system, with enforcement through the WTO's dispute mechanism. TRIPS is administered by the WTO, not WIPO, but the two organisations cooperate closely, and TRIPS incorporates substantive provisions of the Paris and Berne Conventions. Second, the WIPO Copyright Treaty (WCT) and the WIPO Performances and Phonograms Treaty (WPPT), both adopted in 1996 and known as the "Internet Treaties," updated copyright for the digital environment, addressing matters such as the reproduction right in digital form and protections for technological measures.

The digital turn also produced the legal building blocks now invoked in the AI debate. The concept of text and data mining, and the exceptions that permit it, emerged from the data-driven research and analytics of the 2000s and 2010s — long before generative AI — and are now the principal legal hook for AI training in Europe. The fair-use doctrine, similarly, was stretched and tested through cases on search engines, book digitisation (notably the Google Books litigation, which found mass scanning to create a searchable index to be transformative fair use), and other large-scale copying. These precedents form the doctrinal terrain on which today's AI-training cases are being fought.

6.4 WIPO's engagement begins (2019–2021)

WIPO recognised the coming collision early. In 2019 it published WIPO Technology Trends: Artificial Intelligence, a landmark patent-landscape study mapping the explosive growth of AI invention, and convened the first session of what became the WIPO Conversation on Intellectual Property and Frontier Technologies. In 2019–2020 WIPO circulated an Issues Paper on Intellectual Property Policy and Artificial Intelligence, later revised after extensive public comment, which set out — without resolving — the core questions this committee now faces: whether AI-generated output should be protected and by whom; whether AI can or should be named an inventor; how training data implicates copyright; and how the development gap should be handled. The Issues Paper remains a useful map of the problem space and a model of how a norm-setting body frames a contested question without prejudging it.

Through this period the debate was still largely theoretical. The systems of 2019 could generate impressive narrow outputs, but the prospect of a general-purpose technology producing publishable text, gallery-quality images, and working code on demand — and being trained on essentially the entire digital commons to do so — had not yet materialised at scale.

6.5 The generative inflection (2022–present)

That changed from late 2022, when large language models and diffusion-based image generators reached the mass market. Within a remarkably short period, every abstract question in WIPO's Issues Paper became a live, litigated, and legislated controversy. Patent offices confronted the DABUS applications working their way through their systems. Copyright offices confronted registration applications for AI-assisted works. Courts confronted lawsuits from authors, artists, news organisations, and photo agencies alleging mass infringement in training. Legislatures, above all in the European Union, moved to regulate. And WIPO's Conversation series, which had begun as a forward-looking dialogue, became the central global venue for making sense of a landscape changing month to month — culminating in its 2025 session on synthetic media and the 2026 launch of the AI Infrastructure Interchange. The remainder of this guide examines that landscape in detail.

---

7. How Generative AI Stresses the IP System

It is useful to divide the agenda into distinct problems, because almost every controversy delegates will debate sits within one of them. The "output problem" concerns whether and to whom IP rights attach to AI-generated material, and divides naturally into copyright (authorship) and patent (inventorship). The "input problem" concerns whether training on protected works infringes the rights of their creators. A separate but related cluster concerns infringing outputs and who bears liability for them. Synthetic media and likeness form a fourth cluster, and the development and traditional-knowledge dimension cuts across all of them.

7.1 The output problem (I): authorship and copyrightability

The dominant answer so far is that purely AI-generated content is not protected by copyright, and that a human who merely prompts a system generally cannot claim authorship of the raw output. The clearest statement comes from the United States. Stephen Thaler sought to register a visual work, A Recent Entrance to Paradise, naming his AI system the "Creativity Machine" as the sole author. The Copyright Office refused, and on 18 March 2025 the D.C. Circuit Court of Appeals affirmed in Thaler v. Perlmutter, holding that the Copyright Act requires a work to be authored, in the first instance, by a human being (Thaler v. Perlmutter, D.C. Cir. 2025). On 2 March 2026 the U.S. Supreme Court declined to review the decision, leaving it in force (Mayer Brown 2026).

This judicial line is reinforced by the U.S. Copyright Office's policy work. Its Registration Guidance (2023) and the landmark Report on Copyright and Artificial Intelligence, Part 2:

Copyrightability (January 2025) concluded that wholly AI-generated material is not registrable, and — crucially — that prompts alone are generally insufficient to confer authorship, because prompts function as instructions conveying unprotectable ideas and current tools do not give the user enough control over the specific expressive output (US Copyright Office 2025a). The Office's practice, developed through cases such as the graphic novel Zarya of the Dawn and the AI-generated image Théâtre D'Opéra Spatial, distinguishes between protectable human-authored elements — text, selection, arrangement, meaningful editing — and unprotectable AI-generated elements within the same work. Where exactly the line of "sufficient human contribution" falls remains the live, fact-specific question, and the Office expressly anticipated that courts would refine it.

A crucial international contrast prevents this from being a settled global rule. China's Beijing Internet Court, in a widely discussed 2023 decision (Li v. Liu), recognised copyright in an AI-generated image on the ground that the user had exercised sufficient intellectual input through detailed prompts, parameter selection, and iterative refinement (Beijing Internet Court 2023). This more "user-as-author" friendly approach, contrasted with the stricter human-authorship gate in the United States, is precisely the kind of cross-border divergence WIPO exists to address. Delegates should treat the U.S. position as influential but not universal.

The stakes of the output question run to the core of the bargain. If AI output is uncopyrightable, it falls into the public domain, freely usable by anyone — which some regard as healthy and others as a disincentive to investment in producing and curating high-quality AI content. The problem becomes especially acute when human and AI contributions are entangled in a single work, as they increasingly are.

7.2 The output problem (II): inventorship and patents

On the patent side, the global answer has been even more consistent, and again it traces to one man's worldwide campaign. Through the "Artificial Inventor Project," Stephen Thaler filed patent applications around the world — routed through the PCT — naming his AI system DABUS as the sole inventor of, among other things, a food container and an emergency light beacon. Almost every major patent office and court has rejected AI inventorship under existing law:

In the United Kingdom, the Supreme Court held in December 2023 that an inventor must be a natural person, and the applications were deemed withdrawn; the High Court confirmed in 2025 (Thaler v. Comptroller-General, [2025] EWHC 2202 (Ch)) that Thaler could not rescue the applications by belatedly naming himself, having always maintained that DABUS was the true inventor. In the United States, Thaler v. Vidal (Fed. Cir. 2022) held that an inventor must be a human, and the Supreme Court declined review. The European Patent Office, Germany,

Australia, Japan (IP High Court, January 2025), Canada (Patent Appeal Board, June 2025), and Israel (district court, December 2025) have all reached the same conclusion under their respective laws. South Africa granted a DABUS patent, but conducts only a formalities examination without substantive scrutiny of inventorship, and is treated as an outlier rather than an endorsement.

The more consequential development is the emerging treatment of AI-assisted invention. The settled question — can a machine be a named inventor? — has effectively been answered (no). The open question is how much human involvement is required when AI contributes to an invention. The U.S. Patent and Trademark Office's February 2024 guidance holds that AI-assisted inventions remain patentable provided a natural person made a significant contribution to the claimed invention, applying the established Pannu factors for joint inventorship (USPTO 2024). In June 2025 the Swiss Federal Administrative Court took a notably accommodating line, indicating that a natural person can qualify as a co-inventor where they meaningfully influenced the AI by contributing to the data processing that led to the invention (Swiss Federal Administrative Court, B-2532/2024). Germany has accepted naming the human applicant as inventor alongside a statement that the AI was involved.

The emerging consensus is therefore twofold: a machine cannot be a named inventor, but using a machine does not by itself defeat patentability, provided a human's contribution is genuine. The unresolved danger is a coverage gap — inventions generated with so little human input that no person qualifies as inventor may fall outside the patent system entirely, potentially discouraging investment in highly autonomous research and development.

7.3 The input problem: training data and infringement

This is the most economically explosive issue and likely the centre of gravity for committee debate. Generative models learn from enormous datasets scraped from the open web, digitised books, image libraries, and code repositories — overwhelmingly without licences. Rightsholders characterise this as mass infringement on an unprecedented scale; developers characterise it as transformative learning that no licensing market could realistically supply at the required volume and diversity. Courts have split, and revealingly, even courts in the same district have reasoned differently.

The first U.S. ruling to address fair use in an AI-training context went against the developer. In Thomson Reuters v. Ross Intelligence (D. Del., 11 February 2025), Judge Bibas held that Ross's use of Westlaw's editorial headnotes to train a competing legal-research tool was not fair use, finding that the first factor (purpose and character — the use was commercial and not transformative, serving the same function as the original) and the fourth factor (market effect — the use threatened Thomson Reuters's actual and potential markets, including a potential market

for AI training data) both favoured the rightsholder (Thomson Reuters v. Ross, D. Del. 2025). The court was careful to note that Ross's tool was not generative AI, but the decision became among the most-cited in subsequent litigation and the ruling was certified for interlocutory appeal.

Two June 2025 decisions from the Northern District of California then reached the opposite headline result for generative models, by different routes. In Bartz v. Anthropic, Judge Alsup held that training large language models on lawfully acquired books was "exceedingly transformative" and therefore fair use, but sharply distinguished the separate act of building and retaining a central library of pirated copies, which he found was not protected (Bartz v. Anthropic, N.D. Cal. 2025). Anthropic subsequently agreed to a settlement of roughly USD 1.5 billion covering around 500,000 works — described as the largest copyright settlement in U.S. history (Wolters Kluwer 2025). In Kadrey v. Meta, Judge Chhabria also found training to be fair use, but largely because the authors failed to prove market harm; he warned pointedly that a better-evidenced theory of market dilution could change the outcome in future cases (Kadrey v. Meta, N.D. Cal. 2025). The two California judges thus agreed on the result but diverged on reasoning — Alsup focusing on how the data was acquired, Chhabria on the integrated nature of the training process and the centrality of market harm.

The U.S. Copyright Office weighed in with the pre-publication Part 3: Generative AI Training (May 2025). Its central message is that AI training is not categorically fair use; it is a case-by-case, fact-specific inquiry. Highly transformative, general-purpose uses tend to favour fair use, but training a model to generate expressive content that competes with the originals — especially commercially, and especially through pirated access — goes beyond established fair-use boundaries. The Office argued that licensing markets are emerging and viable and favoured voluntary licensing over new legislation (US Copyright Office 2025b). The report's release was accompanied by political turbulence — the Register of Copyrights was dismissed the day after publication — underscoring how contested this terrain is.

7.4 Fair use versus text-and-data-mining exceptions

A structural point for delegates: the world is not merely disagreeing about how to apply one rule; it is operating under fundamentally different legal architectures.

The United States uses the open-ended, judge-made fair-use doctrine, which gives courts flexibility but produces uncertainty and case-by-case outcomes. The European Union, by contrast, does not use fair use at all. It relies on the closed list of exceptions in the 2019 Copyright in the Digital Single Market (CDSM) Directive, two of which cover TDM: Article 3 provides a mandatory exception for scientific research by research and cultural-heritage institutions, with no opt-out; Article 4 provides a broader exception for general and commercial

use, but subject to a rightsholder opt-out — meaning rightsholders can reserve their rights, ideally through machine-readable means, and thereby exclude their works from training.

Layered on top, the EU AI Act (Regulation 2024/1689) imposes obligations on providers of general-purpose AI models: they must adopt a policy to respect rights reservations and opt-outs under the CDSM Directive, and must publish a "sufficiently detailed summary" of the content used for training, following a template issued by the Commission. These obligations began to apply from August 2025, with models already on the market given until August 2027 to comply, and are elaborated in an accompanying General-Purpose AI Code of Practice (European Commission 2025a; WilmerHale 2025). The transparency requirement has been criticised as inadequate — the European Parliament in early 2026 called the summary obligation insufficient because it does not let an individual creator verify whether their specific work was used (Covington 2026) — and the EU has continued consulting on how to make opt-outs and transparency technically workable.

Other jurisdictions sit elsewhere on the spectrum. Japan has a notably permissive exception (Article 30-4 of its Copyright Act) for "information analysis," widely read as friendly to AI training and adopted as a deliberate pro-innovation stance. The United Kingdom consulted in 2025 on an EU-style opt-out TDM exception paired with greater transparency, but met fierce resistance from its creative industries and has not settled on a direction.

The takeaway for WIPO is stark: the world is fragmenting into incompatible regimes — open-ended fair use, opt-out TDM, permissive analysis exceptions, and outright uncertainty — which is exactly the cross-border friction a global norm-setting body might try to bridge.

7.5 Output infringement and the liability question

Distinct from training is the question of infringing outputs — when a model produces something that reproduces or closely imitates a protected work — and who is liable for it. The training cases stressed that even where training is fair use, anyone who publishes an infringing output remains liable. This question is now being litigated directly. In June 2025, Disney and Universal jointly sued Midjourney in California, alleging that its image generator functions as an engine for producing unauthorised images of iconic protected characters, and Warner Bros. later brought a similar suit (Justia 2025; Press Democrat 2025). These cases target outputs rather than (or in addition to) training, and they raise the question of how much responsibility the model provider bears for what users generate.

The UK's first major ruling illustrates the complexity. In Getty Images v. Stability AI ([2025] EWHC 2863 (Ch), 4 November 2025), Getty largely lost: it had narrowed or abandoned its principal training-copyright claims, partly because the training occurred outside the UK, and the

High Court dismissed its secondary-infringement claim while granting only a narrow trademark win where Stability's outputs reproduced Getty's watermarks (Getty v. Stability AI, EWHC 2025; Press Democrat 2025). The case highlighted a recurring theme: the territorial difficulty of pursuing developers whose training happened in another jurisdiction, and the resulting calls — including from Getty itself — for stronger transparency rules so that creators can even identify whether and where their works were used.

7.6 Synthetic media, deepfakes, and digital replicas

WIPO's twelfth Conversation (October 2025) centred on synthetic media: AI-generated images, audio, video, and text. The promise is the democratisation of creativity; the peril is hyper-realistic deepfakes and the unauthorised cloning of real people's voices and likenesses (WIPO 2025b). Existing IP law fits awkwardly here, because a voice or a face is not always a "work" in the copyright sense. Countries are therefore reaching for a patchwork of tools: personality and publicity rights; the U.S. Copyright Office's Part 1 report on digital replicas (2024); consumer-protection and data-protection law (including the EU's GDPR principles of fairness and accountability); contract; and technical measures such as watermarking, labelling, and provenance and traceability mechanisms. Courts in some jurisdictions, including India, have moved to protect performers from voice cloning without consent. Some experts argue the gaps justify an entirely new category of right. Because likeness harms are inherently cross-border — a deepfake created in one country harms a person in another — this is fertile ground for a WIPO framework.

7.7 The development and traditional-knowledge dimension

WIPO is not a club of wealthy, AI-producing states, and a serious resolution will take the development gap seriously. Several concerns recur for low- and middle-income economies. The first is concentration: the compute, capital, and data needed to build frontier models sit with a few firms and countries, and rules written for that frontier may entrench it. The second is asymmetric burden: heavy transparency and licensing regimes may be navigable for large incumbents but crushing for smaller developers and start-ups in the Global South, risking a future in which developing economies are permanent consumers, not producers, of AI. The third is extraction: globally scraped datasets may include the cultural output and traditional knowledge of communities who gave no consent and see no benefit. This last point connects directly to WIPO's most recent treaty achievement — the WIPO Treaty on Intellectual Property, Genetic Resources and Associated Traditional Knowledge, adopted in May 2024, which introduced a disclosure requirement for patents based on genetic resources and associated traditional knowledge (WIPO 2024b). The fourth concern is access and capacity: developing-country delegates frequently prioritise technology transfer, capacity-building, and flexibilities (for

research and education) and seek assurance that new norms will not raise the cost of accessing AI tools.

The framing lesson is that a resolution that only protects rightsholders, or only frees developers, will struggle to win a broad coalition. The development bloc is where balanced language is forged.

---

8. Doctrinal Tools and Trade-offs

Having mapped the problems, it is worth organising the levers available — to WIPO and to its members — around the trade-offs each entails. Delegates should treat these as a menu, not a hierarchy.

Defining the human-contribution threshold. The most consequential single choice on the output side is where to draw the line of "sufficient human authorship" (and the parallel "significant human contribution" for inventorship). A high threshold protects the principle that IP rewards human creativity but leaves vast quantities of AI-assisted work unprotected and in legal limbo; a low threshold extends protection more generously but risks granting monopolies over output that required little human creativity. WIPO could articulate shared principles or factors without dictating a single bright line.

Choosing the training default: opt-in, opt-out, or case-by-case. Requiring permission before training (opt-in) maximally protects creators but may be impractical at frontier scale and would advantage incumbents who can afford licences. Permitting training unless rights are reserved (opt-out, the EU model) balances the interests but depends on workable machine-readable standards and places the burden on creators to act. Case-by-case fair-use-style assessment preserves flexibility at the cost of certainty. A global baseline that accommodates national variation is the realistic WIPO ambition.

Designing compensation mechanisms. If training is to be permitted, creators may still be owed payment. Options include voluntary collective licensing (low-friction but dependent on market formation), statutory or compulsory licensing (certain but administratively heavy and politically fraught), and levy schemes. Each trades off the interests of creators, developers, and the feasibility of administration.

Transparency and its limits. Disclosure of training-data sources and respect for machine-readable opt-outs empower rightsholders to enforce their rights, but must be balanced against trade-secret

protection, the genuine technical difficulty of cataloguing a trillion-token dataset, and the burden on smaller developers. The European debate — between an industry that finds detailed disclosure onerous and creators who find summary disclosure useless — illustrates the trade-off precisely.

Addressing synthetic media. Likeness and voice can be protected through extension of IP, through new sui generis rights, through coordination with personality and data-protection regimes, or through provenance and labelling standards. The choice between a rights-based and a standards-based approach has different enforcement and innovation implications.

Embedding the development dimension. Capacity-building, model exceptions for research and education, technology transfer, and safeguards against extractive data practices are not add-ons but the price of a broad coalition. A framework silent on development will not command consensus.

WIPO's own role. Finally, the committee must decide the form of its action: non-binding guiding principles; a standing observatory or study mechanism; model provisions; a harmonised technical standard for rights reservation and transparency; or a mandate to begin negotiating a binding instrument. Ambition trades off against achievability, and the higher the ambition, the more the room will fracture.

---

9. Major Stakeholders and Bloc Positions

WIPO operates by consensus, and its alignments do not map neatly onto familiar UN groupings. The fault lines here run between producers and consumers of AI, between rightsholder-heavy and technology-heavy economies, and between the developed world and the development bloc.

9.1 AI-producing advanced economies

The United States, the United Kingdom, the European Union and its member states, Japan, the Republic of Korea, Canada, Israel, and Singapore are the principal homes of frontier AI development. They share an interest in legal certainty and in not strangling a strategically important industry, but they diverge sharply on method — the United States favouring fair use and market solutions, the EU favouring opt-out and mandatory transparency, Japan favouring permissive exceptions. They are not a bloc; their internal disagreement is itself a central dynamic of the committee.

9.2 China and major emerging technology powers

China is simultaneously a frontier AI producer and a voice for the developing world, with a distinctive, more permissive output-side jurisprudence and a strong emphasis on sovereignty and non-imposition of external conditions. It is frequently the swing weight between the Western blocs and the Global South. Other large emerging economies with growing technology sectors — India, Brazil, Indonesia, the UAE, Saudi Arabia — combine development-bloc instincts with ambitions to build domestic AI capability, giving them an interest in flexibilities and capacity-building as well as in not being permanently locked out.

9.3 The creative-industries and rightsholder coalition

Cutting across geography is the coalition of rightsholders: publishers, the recorded-music and film industries, news organisations, photo agencies (Getty being the most litigious), visual artists, and the collecting societies and author guilds that represent them. Their consistent demands are consent, compensation, attribution, and transparency, and they have been the driving force behind both litigation and the EU's transparency rules. States with large creative sectors — across Europe, but also in places such as Nigeria (Nollywood), India (its film and music industries), and the Republic of Korea — feel this pressure acutely.

9.4 The development bloc and the Global South

The African Group, the least-developed countries, and many states across Asia, Latin America, and the Pacific prioritise access, development, capacity-building, technology transfer, flexibilities for research and education, and protection of traditional knowledge and cultural expressions. For these delegates the central question is distributional: will new norms widen or narrow the global gap? They are natural champions of balanced language and of routing concerns through the WIPO Development Agenda.

9.5 Non-state actors relevant to debate

Although WIPO's membership is governmental, several non-state actors are routinely referenced: the AI-developer industry and its associations; rights-holder collecting societies and author and artist guilds; standards bodies working on content provenance and watermarking; and civil-society and academic voices on access, privacy, and human rights. Delegates should not propose non-state actors as parties to a resolution, but may cite their analysis and may propose mechanisms (consultations, multistakeholder platforms) that involve them — building on the model of the WIPO Conversation itself.

---

10. Past International Action

International action on AI and IP has moved, in a few short years, from theoretical dialogue to a dense and uncoordinated thicket of national rules and litigation. WIPO itself led the early dialogue: the 2019 Technology Trends report, the launch of the WIPO Conversation that year, and the 2019–2020 Issues Paper on IP Policy and AI framed the questions for the world. WIPO has since sustained that role through successive Conversation sessions — the twelfth, on synthetic media, in October 2025 — and the 2026 launch of the AI Infrastructure Interchange, alongside practical guidance such as Generative AI: Navigating Intellectual Property.

National and regional action has been more consequential but also more fragmented. The United States produced the Copyright Office's three-part AI report (2024–2025) and the USPTO's 2024 inventorship guidance, against a backdrop of major litigation (Thomson Reuters v. Ross; Bartz v. Anthropic and its record settlement; Kadrey v. Meta; The New York Times v. OpenAI; Disney and Universal v. Midjourney). The European Union enacted the CDSM Directive (2019) and the AI Act (2024), the latter introducing the GPAI transparency-summary and opt-out-respect obligations that began applying in 2025. The United Kingdom ran a contentious 2025 consultation and saw the first UK training-and-trademark ruling in Getty v. Stability AI. Japan maintained its permissive Article 30-4 exception. China issued generative-AI and deep-synthesis regulations and developed its more permissive output-side case law.

On the development side, the most significant multilateral achievement has been WIPO's own GRATK Treaty (2024), which — though not about AI — established a disclosure principle for inventions drawing on genetic resources and traditional knowledge that is directly relevant to concerns about extraction. Across all of this, the recurring lesson is that the international system has produced energetic but incompatible responses, leaving cross-border gaps, forum-shopping risks, and a widening divide between the jurisdictions that produce AI and those that merely receive it — precisely the coordination problem WIPO exists to address.

---

11. Current Issues and Points of Contention

Several fault lines will structure debate.

The first is the threshold of human contribution. There is broad but not universal agreement that machines cannot themselves be authors or inventors; the genuine disagreement is over how much human involvement earns protection for AI-assisted work, and whether that line should be harmonised internationally or left to national discretion. A high threshold protects the

human-creativity principle but leaves enormous quantities of work unprotected; a low one risks cheap monopolies.

The second is the training default and the input problem. States and stakeholders divide between an opt-in (permission-first) model favoured by many rightsholders, an opt-out model favoured by the EU and by some developers, and an open-ended fair-use approach favoured in the United States. Beneath this sits the unresolved question of whether and how creators should be compensated when their works train models, and whether voluntary licensing markets can form fast enough to matter.

The third is transparency versus feasibility and trade secrets. Creators — and even well-resourced rightsholders such as Getty — argue they cannot enforce rights they cannot detect, and demand disclosure of training data. Developers respond that detailed disclosure is technically burdensome, commercially sensitive, and especially onerous for smaller players. The European fight over whether the AI Act's summary is meaningful is a preview of this debate at WIPO.

The fourth is synthetic media and likeness, where the contention is less about whether protection is needed than about its form: extension of existing IP, a new sui generis right, reliance on personality and data-protection law, or technical provenance standards.

The fifth, and the one most likely to determine whether any resolution achieves consensus, is the development divide. Developing-country delegates will resist frameworks that entrench the concentration of AI capability or that burden their nascent industries, and will insist on capacity-building, flexibilities, technology transfer, and protection of traditional knowledge. Advanced economies will resist mandates they regard as costly or as threats to competitiveness. Bridging this divide is the central diplomatic challenge of the committee — and, as in the broader multilateral system, consensus has become harder to achieve precisely as the stakes have risen.

---

12. Questions to be Answered

1. Should WIPO articulate a shared principle on human authorship and human inventorship, and if so, how should it define the threshold of human contribution that earns protection for AI-assisted works?
2. How should member states address the coverage gap created by inventions or works generated with so little human input that no person qualifies as author or inventor?

3. What should the default rule be for the use of copyrighted works in AI training — permission-first (opt-in), permitted-unless-reserved (opt-out), or case-by-case assessment — and can a global baseline coexist with national variation?
4. Should WIPO encourage voluntary collective licensing, statutory or compulsory licensing, or levy mechanisms to compensate creators whose works are used to train models?
5. Should there be an international standard for disclosing training-data sources and for machine-readable rights reservation, and how should it be balanced against trade secrets, technical feasibility, and the burden on smaller developers?
6. How should synthetic media, deepfakes, and digital replicas be addressed — through extension of IP, a new sui generis right, coordination with personality and data-protection regimes, or provenance and labelling standards?
7. How can any new framework guarantee capacity-building, technology transfer, flexibilities for research and education, and protection of traditional knowledge, so that it does not widen the global divide?
8. How should liability for infringing AI outputs be allocated between developers and users, and what does that imply for cross-border enforcement?
9. What is the appropriate form of WIPO action — non-binding principles, a standing observatory or study mandate, model provisions, a harmonised technical standard, tasking the standing committees, or the launch of treaty negotiations?
10. How can WIPO improve coordination with the WTO (TRIPS), UNESCO, and other bodies to reduce fragmentation and forum-shopping?

---

13. Further Reading and Bibliography

WIPO — Intellectual Property and Frontier Technologies; the WIPO Conversation; AI Infrastructure Interchange (AIII): https://www.wipo.int/en/web/frontier-technologies

WIPO — Generative AI: Navigating Intellectual Property: https://www.wipo.int/en/web/frontier-technologies

WIPO — Treaty on Intellectual Property, Genetic Resources and Associated Traditional Knowledge (2024): https://www.wipo.int/meetings/en/2024/wipo_ip_gr_tk_dc.html

U.S. Copyright Office — Copyright and Artificial Intelligence, Parts 1–3 (2024–2025): https://www.copyright.gov/ai/

Thaler v. Perlmutter, D.C. Circuit (18 March 2025); cert. denied, U.S. Supreme Court (2 March 2026): https://media.cadc.uscourts.gov/opinions/docs/2025/03/23-5233.pdf

Thaler v. Vidal, U.S. Court of Appeals for the Federal Circuit (2022); UK Supreme Court, Thaler v. Comptroller (2023); UK High Court, [2025] EWHC 2202 (Ch)

USPTO — Inventorship Guidance for AI-Assisted Inventions (February 2024): https://www.federalregister.gov/documents/2024/02/13/2024-02623/inventorship-guidance-for-ai-assisted-inventions

Thomson Reuters v. Ross Intelligence, D. Del. (11 February 2025): https://www.jenner.com/en/news-insights/client-alerts/court-decides-that-use-of-copyrighted-works-in-ai-training-is-not-fair-use-thomson-reuters-enterprise-centre-gmbh-v-ross-intelligence-inc

Bartz v. Anthropic, N.D. Cal. (June 2025) and the 2025 settlement: https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/

Kadrey v. Meta Platforms, N.D. Cal. (June 2025): https://www.jw.com/news/insights-kadrey-meta-bartz-anthropic-ai-copyright/

Getty Images v. Stability AI, [2025] EWHC 2863 (Ch) (4 November 2025): https://www.pressdemocrat.com/2025/11/04/getty-images-ai-lawsuit-britain/

Disney and Universal v. Midjourney (C.D. Cal., filed June 2025): https://news.justia.com/disney-and-universal-file-copyright-lawsuit-against-midjourney-over-ai-generated-content

EU — Directive 2019/790 (CDSM, Articles 3–4); EU AI Act (Regulation 2024/1689) and the GPAI Code of Practice and training-data summary template: https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/european-commission-releases-mandatory-template-for-public-disclosure-of-ai-training-data

The Berne Convention, the Paris Convention, the Patent Cooperation Treaty, and the WIPO Copyright Treaty (treaty texts): https://www.wipo.int/treaties/en/