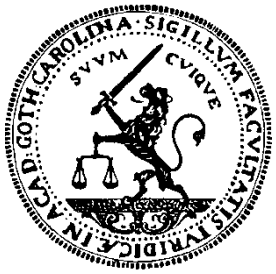

FACULTY OF LAW

LUND UNIVERSITY

Olivia Myrén

Decoding developer authorship

Framing authorship in European Union copyright law in the age of generative AI

JAEM03 Master Thesis

European Business Law

30 higher education credits

Supervisor: Aurelija Lukoševičienė

Term of graduation: Spring 2025

Table of contents

Summary

This thesis examines whether the developers of generative artificial intelligence (GENAI) systems can, and should, be recognised as authors under EU copyright framework for the images generated by their systems. Technological advances over the past decade have redefined traditional conceptions of creativity and copyright in ways that current legislation never anticipated. This study addresses a significant research gap, as previous studies have primarily focused on recognising the systems and their users as authors, while the copyright interests of developers have received limited attention.

The research is presented in three parts. The first part is technical in nature and examines the choices that developers can make during the development process of image generation systems. The second part outlines the legal norms that constitute the EU framework for authorship, with a primary focus on the InfoSoc Directive and the case law of the European Court of Justice. Subsequently, these legal norms are applied to a developer's scenario. The aim is to establish whether the development process enables developers to exercise the necessary free and creative choices to be recognised as authors of the images produced by their generative systems with user input. This systematic review shows that, although many choices are technical in nature, developers often exercise significant creative discretion. However, the system's design and the absence of direct causality between the developer's input and the final image can complicate their copyright claims. The analysis also includes the impact of user autonomy on the assertion of developers' free and creative choices.

The third part discusses whether classical justification theories such as utilitarianism, labour theory and personality theory support the recognition of developers as authors. To answer this question, it examines under what conditions each theory justifies the rights that EU copyright law entails, and to what extent the rights of developers fulfil these conditions. The basis of the investigation reveals that, from a utilitarian perspective, recognising developers' rights can be justified when the economic incentives from GENAI systems are insufficient to stimulate investment, provided that the scope of protection remains proportional. Furthermore, labour theory suggests that significant work effort does not necessarily entitle one to full copyright protection, particularly when the indispensable contribution of the user is considered during the generation process. Lastly, it is concluded that personality theory only advocates recognition when the developer's creative idea is clearly defined through the system's fine-tuning and filtering processes.

Sammanfattning

Denna uppsats undersöker huruvida utvecklare av generativ artificiell intelligens (GENAI) system kan, och bör, erkännas som upphovsmän enligt EU:s upphovsrättsregler för de bilder som genereras med deras system. Teknikens framsteg under det senaste årtiondet har omdefinierat de traditionella föreställningarna om kreativitet och upphovsrätt på sätt som den nuvarande lagstiftningen aldrig avsåg. Genom denna ansats fylls en betydande forskningslucka, eftersom tidigare forskning främst har behandlat GENAI systemet alternativt användarna som alternativ för upphovsmannaskap, medan utvecklares upphovsrättsliga intressen hittills har fått begränsad uppmärksamhet.

Undersökningens material redogörs för i tre delar. Den första är av teknisk karaktär och undersöker de val utvecklaren kan göra under utvecklingsprocessen av bildgenereringssystemet. Den andra delen redogör för de normer som utgör EU:s ramverk för upphovsmannaskap, framför allt med fokus på InfoSoc-direktivet samt EU-domstolens praxis. Därefter appliceras de redovisade rättsnormerna på utvecklarens förutsättningar. Syftet är att utreda om utvecklingsprocessen ger utvecklaren utrymme att göra de fria och kreativa val. Dessa val är centrala för den rättsliga utredningen då de ligger till grund för om utvecklaren kan erkännas som upphovsman till de bilder som deras generativa system producerar. Denna systematiska genomgång visar att utvecklare ofta utövar betydande kreativt handlingsutrymme, även om många val är av teknisk karaktär. Samtidigt kan systemets design och avsaknaden av direkt kausalitet mellan utvecklarens insatts och den slutliga bilden försvåra utvecklares upphovsrättskrav. Uppsatsen analyserar också hur användarens autonomi påverkar möjligheten att hävda utvecklarens fria och kreativa val.

Uppsatsens tredje del avhandlar huruvida det finns belägg inom de klassiska rättfärdigandeteorierna utilitarism, arbetsteorin och personalitetsteorin för att erkänna utvecklarna som upphovsmän. För att besvara denna fråga utreds under vilka förutsättningar som respektive teori rättfärdigar de rättigheter som upphovsrätten inom EU medför samt i vilken grad upphovsmannaskap för utvecklaren uppfyller dessa förutsättningar. Underlaget för undersökningen visar att ur ett utilitaristiskt perspektiv kan erkännande av utvecklares rätt motiveras när de ekonomiska incitamenten från GENAI-systemet är otillräckliga för att stimulera investeringar, förutsatt att skyddets omfattning förblir proportionerlig. Vidare antyder arbetsteorin att betydande arbetsinsats inte nödvändigtvis berättigar till fullt upphovsrättsligt skydd, särskilt med beaktande av användarens oumbärliga bidrag vid genereringen. Avslutningsvis dras slutsatsen att personalitetsteorin förespråkar erkännande enbart när utvecklarens kreativa idé tydligt ramats in genom systemets finjustering och filtrering.

Abbreviations

AI	Artificial intelligence
ANN	Artificial neural network
CJEU	Court of Justice of the European Union
DM	Diffusion model
EU	European Union
GAN	Generative adversarial network
GENAI	Generative artificial intelligence
VAE	Variational autoencoder

1 Introduction

1.1 Background

The emergence of generative artificial intelligence (GENAI) has fundamentally disrupted traditional conceptions of creativity and authorship in ways that copyright law was never designed to contemplate. When artificially generated imagery is mentioned, people tend to think of the giant systems that captivated audiences of all ages when they were released to the public. GENAI systems like *DALL-E*[1], *Midjourney*[2] and *Stable Diffusion*[3] seemed to turn anyone who could form a sentence into an expressive artist. While their usage leads to many interesting legal questions, it is not the focal point of this essay. Instead, this essay is inspired by a lesser-known system derived from the project *The Next Rembrandt*. This thought-provoking model was developed to drive the conversation about art and algorithms by creating a picture that could blend in with the works of the late painter Rembrandt. Not because it imitates one of them exactly, but because it produces something that could possibly pass for a Rembrandt original. The system was trained through machine learning to recognise and visualise common denominators in the works to produce something that was both new and reminiscent of what already existed.[4]

Undeniably, a new art form has emerged, offering countless of opportunities for creative minds. Rather than bringing an internal vision to life with a paintbrush, this alternative medium offers a new workflow. By designing and curating datasets, fine-tuning the model, and rigorous testing, a system is created that can generate novel images according to the developer's vision. However, beneath this creative marvel lies a profound legal complexity that strikes at the heart of European Union (EU) copyright doctrine: if copyright protection requires human authorship and originality that reflects an author's personality through *free and creative choices*, what happens when the creative process involves creating a product that depends on the actions of others? Do the images produced truly reflect the developer's personality in this scenario, or are there other factors that justify an authorship claim? Without this recognition, free competition reigns, severely limiting the opportunities to commercially exploit the end products. While systems themselves are often protected, individual artists who wish to produce art using this new medium risk being left unprotected. This thesis is dedicated to artistic developers who bridge technology and art, creating new forms of expression through GENAI.

---

[1] 'DALL·E 3' (*OpenAI*) <https://openai.com/index/dall-e-3/> accessed 11 July 2025.

[2] 'Midjourney' <https://www.midjourney.com/website> accessed 11 July 2025.

[3] 'Stable Diffusion Online' (*Stable Diffusion*) <https://stablediffusionweb.com/> accessed 11 July 2025.

[4] 'The Next Rembrandt: Recreating the Work of a Master with AI' (*Microsoft News Centre Europe*, 13 April 2016) <https://news.microsoft.com/europe/features/next-rembrandt/> accessed 9 July 2025.

1.2 Purpose and research question

This thesis examines the question of developer authorship for images generated by GENAI systems under EU copyright law, addressing fundamental tensions between technological innovation and established copyright doctrine. The research approaches this inquiry through both descriptive legal analysis (de lege lata) and normative theoretical evaluation (de lege ferenda). This method addresses the significant uncertainty regarding both the eligibility and desirability of considering the developer as author of the images generated using the developer's system. This uncertainty generates legal risks for developers investing substantial resources in system development, potentially undermining innovation incentives by creating ambiguity about intellectual property returns, while simultaneously raising concerns about the proper scope of copyright protection in an era of algorithmic creation.

The fundamental legal problem this thesis addresses is that current EU copyright frameworks were designed without GENAI in mind. Unlike traditional creative processes where authors maintain direct control over their expression, GENAI system developers exercise only indirect influence, while users initiate the actual generation process through prompts and parameters. This temporal and causal disconnection between development and generation creates uncertainty about whether and when developers can claim authorship over the generated images, particularly given that developer choices during system design may significantly influence generative outcomes despite the absence of direct creative control during the generation phase.

The significance of this inquiry extends beyond doctrinal analysis to fundamental questions about the justification of intellectual property rights in the digital age. The three principal theories that offer justification for intellectual property rights, utilitarian theory, labour theory, and personality theory, each provide distinct frameworks for evaluating whether copyright protection serves legitimate policy objectives.[5]

[5] Steven Ang, *The Moral Dimensions of Intellectual Property Rights* (Edward Elgar 2013) 62.

The questions of this thesis can be summarised as follows:

1. Under EU copyright law, to what extent can developers of generative AI systems be recognised as authors for images generated by their models, considering existing originality requirements and the concept of 'author's own intellectual creation' established in CJEU jurisprudence?

2. Under EU copyright law, to what extent should developers of generative AI systems be recognised as authors for images generated by their models, based on classical justification theories (including utilitarian, labour and personality-based approaches) and the general principles of EU law?

By addressing these questions, this thesis fills a significant gap in EU copyright scholarship with profound implications for the artificial intelligence (AI) industry, creative professionals, and the broader digital economy. As GENAI systems become increasingly sophisticated and commercially significant, establishing clear legal parameters for copyright protection becomes essential for balancing innovation incentives with the foundational principles of EU copyright law. The research provides frameworks for understanding not only whether developers can claim authorship, but whether such recognition serves the policy objectives underlying EU copyright protection.

1.3 Delimitations

The examination is strictly limited to the question of whether developers can claim authorship of AI-generated images created using their system. This thesis does not address whether the GENAI models themselves can or should obtain authorship status, as this would require engaging with questions of legal personhood that fall outside the established framework.

Furthermore, this thesis does not examine the copyright protectability of AI systems themselves or their constituent components, such as training datasets, nor of generated output. These questions, while relevant to the broader AI copyright landscape, represent distinct analytical challenges that would require separate legal frameworks and methodological approaches. While the concepts of authorship and copyright-protected works are inherently intertwined, the analysis focuses exclusively on the former.

The technical exposition of GENAI systems serves as contextual foundation rather than comprehensive technical guide. The analysis relies on secondary sources and does not involve original empirical research or direct collaboration with computer scientists. Technical processes have been simplified where necessary to maintain focus on legally relevant aspects of developer decision-making.

To avoid overreaching analysis and maintain analytical precision, this thesis is limited to generative models that produce images. While numerous generative model types exist with varied development processes, concentrating on image generation allows for nuanced examination of the specific creative choices that developers make in this domain.

The analysis has two time-specific limitations. First, the analysis of developer contributions extends only to the point of system deployment. While many AI systems undergo continuous monitoring and improvement post-deployment, examining these ongoing processes would introduce complexity that extends beyond the thesis's focused research question. Similarly, the examination excludes detailed analysis of post-processing and output refinement techniques. This delimitation is methodologically justified as post-processing occurs after the core generative model has produced its output, representing a separate analytical layer that would require extensive technical examination of filtering algorithms, quality control mechanisms, and content moderation systems. Second, the analysis will not consider any form of editing or other post-production measures of the generated image. There are too many options and to consider all of them would simply not be possible within the purpose of the thesis while maintaining focus on the core question of the developer.

The thesis provides general analysis without detailed examination of specific developers' creative processes. This approach ensures broader applicability but acknowledges that nuanced aspects of individual production processes may be lost. The development process analysis recognises that developers need not be involved in every step. For instance, pre-processed datasets can be purchased commercially. However, rather than cataloguing all possible variations, the analysis highlights those steps crucial for fulfilling copyright thresholds. Every development process differs, and this thesis provides a general framework for understanding these variations.

Ethical and safety considerations in AI development are deliberately excluded from the developer choice analysis. While these considerations involve significant decision-making by developers, they fall outside copyright law's concern with creative expression. Choices regarding bias mitigation, fairness constraints, and safety alignment, while important for AI governance, are primarily normative and regulatory decisions rather than creative artistic choices that would contribute to authorship claims under the EU copyright framework.

When considering to what extent authorship should be granted based on classical justification theories and general principles of EU law, the scope of this thesis remains strictly confined to copyright law. Alternative protection regimes, such as neighbouring rights, patent, design, trademark, or sui generis database rights, are expressly excluded from analysis. While these frameworks might in some contexts offer forms of recognition or protection for

aspects of AI-driven innovation, their objectives, eligibility criteria, and foundational principles diverge from those in copyright law. Introducing them would risk conflating fundamentally different rationales and expanding the analytical focus beyond the central question of this thesis.

These delimitations collectively ensure that the thesis maintains analytical rigor while addressing the specific legal question of developer authorship within the established framework of EU copyright law. By clearly defining what falls outside the scope of analysis, these boundaries enable focused examination of the central research question while acknowledging the broader context in which it exists.

1.4 Methodology and material

This chapter outlines the methodological foundation and source materials that underpin the analysis. This thesis adopts a dual analytical framework that addresses the inherently different nature of the two research questions posed. The first research question examines the current state of EU copyright regarding developer recognition as authors, a de lege lata inquiry focused on interpreting existing legal frameworks. This question will be answered using legal dogmatic method, and an interdisciplinary approach which will be described below. The second question explores the theoretical foundation that justify exclusive rights within the copyright doctrine, requiring engagement with philosophical justifications for authorship that inform legal policy. This requires a normative analysis that considers the developer's role and the theories in the light of EU law. The following sections will first give an overview of the legal dogmatic method, followed by a discussion regarding the interdisciplinary approach that has been applied.

1.4.1 Legal dogmatic method

The analysis employs the legal dogmatic method as its primary methodological approach. This method serves to interpret and systematise the law through the identification and analysis of legal norms, which are then applied to specific scenarios.[6] The legal dogmatic approach seeks to reconstruct the legal landscape by examining binding sources in accordance with their hierarchical order.

The methodology follows the established framework where legal provisions are combined with descriptions of facts and assessments that, while logically derived, may be contested in ambiguous cases. The goal is not to provide answers to exact scenarios, but general council about the legal landscape.[7]

[6] Mauro Zamboni, 'Rättsdogmatisk Metod [Legal Dogmatic Method]' in Maria Nääv and Mauro Zamboni (eds), *Juridisk metodlära [Legal methodology]* (2nd edn, Studentlitteratur 2018) 21.

[7] Aleksander Peczenik, *Juridikens Teori Och Metod: En Introduktion till Allmän Rättslära [Legal Theory and Method: An Introduction to General Jurisprudence]* (Faculty of Law, Lund University 1995) 34.

Given the complex questions surrounding GENAI output and the EU copyright system, a range of authoritative sources beyond statutory provisions are essential in addition to requiring the application of various interpretative methods to reach sound legal conclusions. [8]

The research includes treatment of legal sources in accordance with established EU legal methodology. This hierarchy places primary EU law at the apex, followed by secondary legislation, then court decisions and preparatory work, with legal doctrine providing additional analytical insight without binding effect. [9] CJEU has historically relied on international instruments to interpret EU copyright law, particularly the *Berne Convention for the Protection of Literary and Artistic Works* (Berne convention), reflecting the significant influence of these foundational agreements.[10] The analysis therefore incorporates relevant international instrument as directed by CJEU precedent, alongside the previously mentioned sources. [11]

The interpretation of EU law within this thesis employs multiple recognised methods of legal analysis. First, textual interpretation, examines the linguistic meaning of statutory provisions, drawing upon legal definitions, dictionaries, and ordinary language usage. Second, systematic interpretation considering legal provisions within their broader context, examining their relationship to other parts of relevant legal texts, the structural organisation of the law, and established legal-dogmatic theories. Third, teleological interpretation that adopts a goal-oriented examination that considers the objectives underlying the legislation. Fourth, analogous interpretation which compares provisions with closely related law, precedent, or unwritten norms, following the principle of equal treatment.[12]

1.4.2 Interdisciplinary approach

The technical nature of GENAI necessitates investigation into computer science literature to provide a foundation for legal evaluation. An understanding of this technical context becomes crucial for the formation of informed legal analysis. Misrepresentation of technological processes would compromise the quality and reliability of the thesis's legal findings.[13]

---

[8] ibid 46.

[9] Jane Reichel, 'EU-Rättslig Metod [EU Legal Method]' in Maria Nääv and Mauro Zamboni (eds), *Juridisk metodlära [Legal methodology]* (2nd edn, Studentlitteratur 2018) 126.

[10] Berne Convention for the Protection of Literary and Artistic Works (adopted 9 September 1886, as revised at Paris 24 July 1971, and amended 28 September 1979) (Berne convention).

[11] Eleonora Rosati, *Copyright and the Court of Justice of the European Union* (First edition, Oxford University Press 2019) 54.

[12] Peczenik (n 7) 50–59.

[13] Claire Largan and Theresa Morris, *Qualitative Secondary Research: A Step-By-Step Guide* (SAGE Publications Ltd 2019).

The integration of technical and legal analysis presents unique challenges that differentiate this research from traditional legal scholarship. Unlike established legal doctrine with predictable hierarchical structures, technical literature encompasses diverse publication types of varying levels of peer review, and in the case of GENAI, rapidly evolving knowledge domains. To ensure qualitative material collection, the selection has followed structured criteria to ensure reliability and representativeness. Primary focus remains on academic and research sources, with original technical materials interpreted through secondary scholarly analysis where appropriate.

The research assesses credibility, accuracy, reasonableness, and support according to the CARS evaluation framework to ensure source quality. Credibility evaluation in this case has meant examining authoritative affiliations, prioritising peer-reviewed materials due to its gold standard of knowledge production while avoiding purely commercial sources unless providing objective data. Accuracy considerations account for the rapid pace of technological advancement, ensuring currency of technical information. Reasonableness assessment evaluates whether data is representative, meaningful, and comprehensible, with careful attention to context and consistency. Support requirements focus on sources with transparent documentation and acknowledged references.[14]

1.5 Terminology

Evidently, the thesis centres on the relationship between developers of GENAI systems and the images produced by their models. For analytical clarity, this thesis employs *developer* as a collective term encompassing all individuals who contribute creative and technical decisions during system development from initial problem definition through final deployment but are crucial not involved in the generative process. In practice, GENAI development involves multidisciplinary teams including data scientists, machine learning engineers, architects, and product designers.[15] While these roles may be distributed across multiple persons or organisations, the term *developer* captures the aggregate of development-phase contributions that may bear legal significance for authorship claims. This approach enables systematic analysis of which development actions and decisions should be considered when evaluating authorship eligibility, regardless of how responsibilities are allocated within development teams. In contrast, the *user* refers specifically to individuals who interact with deployed GENAI systems during the generation phase, providing prompts or other inputs that trigger image creation. This distinction proves crucial for understanding the temporal and causal relationships between different types of creative contributions.

[14] ibid.

[15] Daswin De Silva and Damminda Alahakoon, 'An Artificial Intelligence Life Cycle: From Conception to Production' (2022) 3 Patterns 100489.

1.6 Research landscape

1.6.1 To what extent can developers of generative AI systems be recognised as copyright holders?

Within the EU, there has been cases in national courts that has questioned whether outputs generated with GENAI are eligible for copyright. The case law points towards a consensus the outputs are generally not in question for copyright. However, a distinct and sufficient human contribution can rebuttal this presumption. The question of whether the developer can contribute in this way has not been object of further discussions.[16]

Academic inquiry into GENAI authorship has been shaped by a predominant focus on two stakeholder categories: the AI system as a potential rights holder and the user as creator.[17] This binary approach has resulted in developer copyright interests receiving only ancillary consideration within broader authorship discussions.[18] The scholarly consensus acknowledges that developer contributions, including training data selection, model architecture design, and system configuration, constitute legally relevant creative decisions, yet this recognition has not translated into systematic analysis of developer authorship eligibility.[19]

The studies thus far have been focused on publicly available broad-domain GENAI models that designed as creative tools for third-party users. This focus has created a methodological gap in examining developer copyright claims. The literature reveals insufficient attention to developer creative contributions throughout the system development process and the complex dynamics between developers and users.[20] This thesis addresses the resulting

[16] 'The Development of Generative Artificial Intelligence from a Copyright Perspective' (European Union Intellectual Property Office (EUIPO) 2024) EUR 31687 EN 69 <https://euipo.europa.eu/ohimportal/en/news/-/action/view/1269306>.

[17] Nicola Lucchi, 'ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems' (2024) 15 European Journal of Risk Regulation 602.

[18] See Alina Škiljić, 'When Art Meets Technology or Vice Versa: Key Challenges at the Crossroads of AI-Generated Artworks and Copyright Law' (2021) 52 IIC - International Review of Intellectual Property and Competition Law 1338; P Bernt Hugenholtz and João Pedro Quintais, 'Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output?' (2021) 52 IIC - International Review of Intellectual Property and Competition Law 1190; Kateryna Militsyna, 'Human Creative Contribution to AI-Based Output – One Just Can('t) Get Enough' (2023) 72 GRUR International 939 (illustrative examples).

[19] Hugenholtz and Quintais (n 18).

[20] European Commission. Directorate General for Communications Networks, Content and Technology., 'Trends and Developments in Artificial Intelligence :Challenges to the Intellectual Property Rights Framework : Final Report.' (Publications Office 2020) <https://data.europa.eu/doi/10.2759/683128> accessed 7 January 2025; Ana Ramalho, *Intellectual Property Protection for AI-Generated Creations: Europe, the United States, Australia and Japan* (1st edn, Routledge 2021) <https://www.taylorfrancis.com/books/9780367823290> accessed 17 February 2025.

analytical deficit in understanding how developer choices during the entire development process might influence copyright recognition.

Notably, the European Intellectual Property Office (EUIPO) released a comprehensive study on the intersection between copyright and GENAI in May 2025. While this study provides detailed technical, legal and economic analysis of GENAI's functionality and examines the development process, its purpose focuses on copyright-protected works as training data, content generation by AI systems, and broader implications for the copyright ecosystem rather than examining developer authorship eligibility.[21]

1.6.2 To what extent should developers of generative AI systems be recognised as copyright holders?

Despite substantial engagement with classical copyright theories and broad discussion of EU legal principles, the literature largely addresses justifications underpinning copyright protection of outputs rather than focusing on the criteria for recognising authorship itself. [22] Normative exploration of the developer's position is therefore minimal and, where present, tends to be rooted in US utilitarian reasoning rather than EU frameworks and theory.[23]

1.7 Outline

Following this introductory chapter, this thesis is structured into four main chapters that systematically address the research questions through both descriptive legal analysis and normative theoretical examination.

Chapter 2 Generative Artificial Intelligence provides the essential technological foundation necessary for informed legal analysis. This chapter examines the anatomy of GENAI systems, detailing their neural network architecture and the distinction between input layers, hidden layers, and output layers. The development process is dissected through six key phases: problem definition, data collection and preprocessing, architecture selection, model training, evaluation and validation, and fine-tuning. Each phase is analysed with particular attention to the choices available to developers, as these creative and

[21] European Union Intellectual Property Office (EUIPO), 'The Development of Generative Artificial Intelligence from a Copyright Perspective' (European Union Intellectual Property Office 2025) <https://euipo.europa.eu/tunnel-web/secure/webdav/guest/document_library/observatory/documents/reports/2025_GenAI_from_copyright_perspective/2025_GenAI_from_copyright_perspective_FullR_en.pdf>.

[22] See Ryan Abbott and Elizabeth Rothman, 'A Utilitarian Approach to Copyright Law and Generative Artificial Intelligence' in Guillaume Tusseau (ed), *Research Handbook on Law and Utilitarianism* (Edward Elgar Publishing 2024) <https://www.elgaronline.com/view/book/9781789901726/9781789901726.xml> accessed 16 July 2025 (illustrative examples); Jan Smits and Tijn Borghuis, 'Generative AI and Intellectual Property Rights' in Bart Custers and Eduard Fosch-Villaronga (eds), *Law and Artificial Intelligence : Regulating AI and Applying AI in Legal Practice* (TMC Asser Press 2022).

[23] Zachary Catanzaro, 'Beyond Incentives: Copyright in the Age of Algorithmic Production' (2023) 13 NYU Journal of Intellectual Property & Entertainment Law 1.

technical decisions form the basis for potential authorship claims. The generative process is then explored, examining how user inputs, seeds, temperature settings, and sampling parameters influence image generation. This technical exposition culminates in conclusions about the balance of control between developers and users, and the implications of the black box problem for establishing causal relationships between developer choices and final outputs.

Chapter 3 Legal Framework constitutes the core doctrinal analysis of EU copyright law as it applies to the developer scenario. Beginning with the foundational sources of EU copyright law, this chapter examines the general principles of proportionality and high-level protection. The construction of copyright is analysed through its two fundamental criteria: originality and expression. The originality requirement is explored through the established test of *author's own intellectual creation*, examining whether developers can demonstrate the necessary *free and creative choices* and *personal touch* throughout the development process. The expression criterion is applied to determine when developer contributions transition from unprotectable ideas to protectable expressions. Following, joint authorship principles are discussed with focus on how user involvement might affect developer copyright claims. The analysis concludes with an examination of historical challenges posed by photography and computer software, providing instructive teachings that can serve as models for how author's actions are examined. Lastly, the subchapter Developer scenario synthesises the findings from Chapter 2 and Chapter 3, providing answers to the first research question.

Chapter 4 Theoretical Elements of Authorship shifts from descriptive to normative analysis, evaluating whether developer recognition aligns with the fundamental theories justifying copyright protection. Three classical justification theories are examined: utilitarian theory, which asks whether recognising developer authorship maximises social welfare through innovation incentives; labour theory, which considers whether developers' substantial investments of skill, time, and resources warrant proprietary rewards; and personality theory, which examines whether GENAI systems sufficiently embody developers' creative personalities to justify protection. Each theory is applied specifically to the developer scenario, weighing arguments for and against recognition while considering the unique characteristics of GENAI production.

2 Generative artificial intelligence

2.1 Introduction

This initial chapter establishes the technological foundation necessary for the legal analysis of developer authorship image generation using AI. The study begins by providing a brief overview of the historical evolution of image generation. It then proceeds to analyse the architecture and core components of contemporary generative systems, with particular emphasis on image generation. The subsequent subchapter systematically presents the development process, from problem definition through data selection, architecture, training, to fine-tuning. It highlights the specific choices available to developers that may bear legal significance. Following the deployment of the system, the focus shifts to the generative process itself, encompassing the manner in which users interact with the system to generate images, and, of paramount importance, the manner in which the developer design and constrain user agency. The chapter concludes with a summary of the developer's legally relevant choices, thus establishing the foundation for the subsequent legal analysis in *Chapter 3.6*. Throughout, the technical exposition is selectively tailored to foreground those decision points most pertinent to the originality and authorship standards under EU copyright law.

The concept of generating images using AI can be traced back to the early 1970s', marked by the introduction of the system *AARON*. The system produced monochrome illustrations by employing symbolic, rule-based programming. This approach relied on symbolic manipulation, operated under conditions that were not entirely predetermined. Due to the rule-based random generation process, the images demonstrated variability that is reminiscent of the contemporary GENAI images.[24] The early 2020's marked a significant uptick in the use of AI, particularly in the field on GENAI. This surge was facilitated by decades of technological advancements, particularly the increased access to data brought by the Internet. [25] These developments enabled the creation of models that are now employed globally. Prominent examples include DALL-E, Midjourney, and Stable Diffusion, which have substantially broadened the range of creative possibilities for visual artists.

GENAI is a subfield of AI that focuses on the development of systems capable of generating new content, including images, video, and music. [26] These systems are trained on extensive datasets and learn statistical patterns, thereby enabling them to create outputs with similar characteristics to their training

[24] Pamela McCorduck, *Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen* (WH Freeman 1991) 7.

[25] European Union Intellectual Property Office (EUIPO) (n 21) 26.

[26] Ajay Bandi, Pydi Venkata Satya Ramesh Adapa and Yudu Eswar Vinay Pratap Kumar Kuchi, 'The Power of Generative AI: A Review of Requirements, Models, Input–Output Formats, Evaluation Metrics, and Challenges' (2023) 15 Future Internet 260, 2.

data.[27] Contemporary generative models encompass large language models (e.g., *ChatGPT* for text), image generators (e.g., *DALL-E*, *Midjourney*, and *Stable Diffusion*), and music generators (e.g., *Suno*). These models can be interacted with in various ways, including text-to-text, text-to-image, or image-to-text. Some advanced multimodal systems can process and generate content using combinations of text, images, and other data types.

Moreover, the scientific discipline of training machines to recognise patterns in data is known as machine learning. [28] Consider an art student learning to recognise different painting styles. Initially, they study hundreds of labelled examples (*this is a Monet*, *this is a Picasso*, *this is a Van Gogh*) building a foundation in pattern recognition across various artistic movements. When the student wishes to specialise in impressionist landscapes, the student can fine-tune their existing knowledge by focusing on specific examples of brushwork, colour palettes, and compositional techniques unique to that style. The more targeted examples they study, the better they become at both recognising and creating works in that particular style, while retaining their broader artistic knowledge. GENAI models work similarly, they pick up patterns from existing pieces (paintings in this case) and learn to produce new designs. Just as the student, the system can be finetuned to produce artwork of a specific style in higher quality. The more specialised data the model observes during fine-tuning, the more accurate and refined its outputs become for that particular creative task.

Furthermore, as with all AI systems, GENAI systems require training in order to accomplish their designated tasks. The common factor between GENAI and other AI systems is pattern recognition, whereby the model adapts to the data to which it is exposed to, leading to improved performance on specific tasks. Machine learning is an expansive field with numerous methodologies, each of which can result in varied outcomes. One such methodology that has gained significant popularity in recent times is the utilisation of artificial neural networks.[29] The networks in question are inspired by the structure of the animal and human brain, consisting of interconnected artificial neurons.[30] Nevertheless, it is imperative to acknowledge that the comparison between the brain and an artificial neural network (ANN) is analogous to the comparison between a bird and an airplane. Despite both aircraft types being capable

[27] European Union Intellectual Property Office (EUIPO) (n 21) 20.

[28] ibid 30.

[29] Michael Negnevitsky, *Artificial Intelligence: A Guide to Intelligent Systems* (2. ed., [Nachdr.], Addison-Wesley 20) 165.

[30] Rudolf Kruse and others, *Computational Intelligence: A Methodological Introduction* (Springer International Publishing 2022) 7 <https://link.springer.com/10.1007/978-3-030-42227-1> accessed 18 February 2025.

of flight, they are fundamentally different in their operational principles. Artificial neurons can be considered a highly simplified abstraction of their biological counterparts.[31]

Moreover, the most versatile models at present are known as foundation models (e.g., GPT-4, which should not be confused with applications such as ChatGPT, although the latter builds upon the former). These models are trained on large-scale, diverse datasets and can serve as a basis for developing other models.[32] With regard to the generation of images, these foundational models have the capacity to produce a broad spectrum of visual art. However, these models are often outperformed by domain-specific models that have been fine-tuned for specific use cases.

In summary, the concept of generative image creation is not a novel one. Nevertheless, the advent of contemporary technological innovations, such as ANNs, has rendered traditional methodologies, including rule-based and symbolic reasoning as exemplified by AARON, almost indistinguishable from traditional AI in the contemporary context. Furthermore, the system requires training in order to perform tasks and can be more or less fine-tuned for specific purposes. In the context of image generation and the examination of this thesis, it is imperative to acknowledge that not all systems capable of generating images possess equivalent capabilities.

2.2 Anatomy

This section provides a systematic breakdown of the core technical components of a GENAI system, thereby establishing a foundation for understanding where and how developer choices arise. This focus will be developed in greater detail in the following sections on the development and generative processes, with particular attention to their legal relevance under EU copyright law.

Starting from the user's access point to the system, which is commonly referred to as the user interface. This viewpoint is customised by the developer, with consideration given to the system's competence and use case.[33] As previously stated, the nature of the input provided by a user to the system is dependent upon the model. For instance, the user can enter text (commonly referred to as a prompt) to when generating the images. In contrast, there are generation systems that do not permit input from the user due to design. For instance, the website *This person does not exist*[34] initiates the generation of a

[31] ibid 10.

[32] Muhammad Awais and others, 'Foundation Models Defining a New Era in Vision: A Survey and Outlook' (2025) 47 IEEE Transactions on Pattern Analysis and Machine Intelligence 2245, 2247.

[33] European Union Intellectual Property Office (EUIPO) (n 21) 31.

[34] 'This Person Does Not Exist' <https://thispersondoesnotexist.com/> accessed 9 July 2025.

human portrait when the user accesses, or refreshes, the website. With the exception of the user interface, the user is not granted any further access to the system.

At its core, the image generator consists of an ANN, which was described briefly in the previous chapter. The ANN comprises artificial neurons, or nodes, that are interconnected by links and systematically arranged in layers. Each connection between nodes is assigned a weight. This weight is initially random but is adjusted throughout the training process to optimise the network's performance. [35] The structure of the network is characterised by several distinct layers. The input layer receives raw data, passing it to one or more hidden layers composed of hidden nodes. These hidden layers are responsible for transforming and abstracting the input through a series of weighted computations. Finally, the output layer produces the network's output. [36] While each connections carry a weight, each node carries a bias value. Both parts play crucial roles in determining the activation of subsequent neurons, enabling the network to progressively identify fundamental features within the input data and combine them into increasingly complex structures. In essence, this process enables the model to discern patterns and generate outputs that mirror the learned data distribution.[37]

Moreover, the architecture of the ANN, specifically the number of layers and number of neurons within each layer, is determined by the developer.[38] In the context of neural networks, the term deep neural network is employed to describe a network that possesses multiple hidden layers. These entities are frequently characterised as *black boxes*, an expression that alludes to the opacity of their internal decision-making processes, from the input to the output layer. This absence of transparency gives rise to considerable challenges with regard to technical interpretability. Indeed, even developers encounter difficulties when attempting to interpret the rationales of their systems.[39]

In summary, the user can only access the system from the user interface that the developer provides, and at times, the input is restricted to only initiation. The system comprises an input layer, an output layer and, in general, at least one hidden layer, located behind the interface. Each layer comprises nodes that are interconnected. It is evident that these parameters are instrumental in the control of the generation process. It is challenging for developers to comprehend the rationale behind their own system's decisions due to the opaque nature of the hidden layers.

[35] Kruse and others (n 31) 40.
[36] ibid.
[37] European Union Intellectual Property Office (EUIPO) (n 21) 26.
[38] Kruse and others (n 31) 107.
[39] ibid 111.

2.3 Development process

This chapter is devoted to a comprehensive analysis of the timeline from conceptualisation to execution. The process incorporates varying degrees of human involvement, a factor that is imperative for the evaluation of authorship status. Following the exposition of each step, a summary is provided that highlights the actions available to the developer in *Chapter 2.3.7*. Specifically, this chapter sequentially examines:

- problem definition
- data collection
- architecture selection
- model training
- evaluation and validation
- fine-tuning.

2.3.1 Problem definition

As previously outlined, an AI model is only capable of performing the tasks for which it has been trained. The problem definition of an image generator defines that this task is, meaning what use will the model have. Consider these two illustrative examples. First, a business whose business model is to monetise the system itself might have the goal to produce a model that users can engage with to produce their own art based on vague ideas. Second, an artist that usually works with photography and has created their own style in this medium. Now they want to portrait unrealistic objects, such as imaginary creatures, in their style. These two scenarios necessitate the implementation of systems that possess distinct capabilities. As will be demonstrated subsequently, the desired objective will have a significant impact on the majority of the steps in the production chain. In addition to defining the objective, it is imperative to consider factors such as hardware and software requirements, system configuration settings, legal implications, and ethical considerations.[40]

2.3.2 Data collection and preprocessing

After defining the problem, the subsequent pivotal steps entail the collection and preprocessing of data. In order to generate new images, a model must be able to learn patterns from a substantial volume of existing data, meaning pieces of art for image-generators. The basis of all systems are foundation models, which require massive and diverse datasets to capture a wide array

[40] Bandi, Adapa and Kuchi (n 26) 8.

of characteristics.[41] In most cases, the training data is consistent with the intended output (e.g. images for image generation). In the context of image generation through text prompts, the utilisation of datasets comprising images accompanied by written captions are commonly used. This facilitates the model's ability to comprehend the correlation between textual descriptions and visual content. To illustrate this point, consider the example of an art student learning to differentiate works by van Gogh, Monet, and Picasso, as outlined in *Chapter 2.1*.[42] Further techniques will be addressed in *Chapter 2.3.4.*

Furthermore, the data is gathered using various methods, often sourcing from multiple origins to ensure breadth and relevance. [43] It is imperative that the dataset represents the patterns and characteristics that are in alignment with the learning objectives that have been established in the problem definition. The selection of collection method depends on multiple factors, notably the type of training data. [44] This choice will also directly affect the quality of the dataset and, by extension, the quality of the images, making it a crucial consideration. [45] There are a number of approaches that may be adopted in order to ascertain the definition of qualitative data. In general, five attributes are considered to be the most important: accuracy, representativeness, consistency, currency and relevance.[46] Prior to its utilisation for training, a dataset undergoes a series of preprocessing steps, including cleaning, with the objective of ensuring optimal performance. This entails elimination of irrelevant data such as duplicates and outliers, editing random errors, inconsistencies, and unwanted noise in images. [47] Depending on the type of training method employed, it may also be necessary to label the data so that the system can learn expected outcomes.[48]

2.3.3 Architecture selection

As with previous stages, the decision of architecture is closely tied to the problem definition, with different models demonstrating varying levels of proficiency in different tasks. Model architecture describes the basic structure of a neural network, specifying how its layers and components are arranged. This foundational design determines how a model processes and generates information, fundamentally shaping the images it generates. [49] Several types of neural architectures are capable of generating images, each employing distinct generative mechanisms. Popular examples include variational autoen-

[41] European Union Intellectual Property Office (EUIPO) (n 21) 32.
[42] ibid 30.
[43] ibid 130.
[44] Bandi, Adapa and Kuchi (n 26) 8.
[45] European Union Intellectual Property Office (EUIPO) (n 21) 30.
[46] ibid 145.
[47] ibid.
[48] ibid 129.
[49] Bandi, Adapa and Kuchi (n 26) 10.

coders (VAEs), generative adversarial networks (GANs), and diffusion models (DMs). All three are probabilistic in nature, meaning they are designed to produce multiple, statistically variable outputs.[50]

The selection of model type exerts a substantial influence on the system's capabilities and the quality of its image output, thus constituting a pivotal step in the developer's process. It is evident that each architectural style proffers a unique set of advantages. Moreover, within each category, there exists a plethora of variants that facilitate further. [51] Furthermore, hybrid models that combine elements from different architectures are becoming increasingly common, allowing developers to leverage the strengths of multiple approaches.[52]

2.3.4 Model training

As previously mentioned, a model without training possesses only random parameters, meaning that an untrained model will inevitably produce images that are random or devoid of meaningful structure. The fundamental objective of the training is to enable the model to discern, internalise, and reproduce patterns from the training data, thereby equipping it to generate plausible results when confronted with novel, unseen inputs. To illustrate this point, if one aims to generate images of dogs, the generative model must be trained on a dataset containing labelled images of dogs. While humans possess the capacity to conceptualise the notion of a dog, the model's understanding is constrained to the statistical patterns inherent in the data with which it is presented.

The training of GENAI models can be achieved through a variety of distinct machine learning methodologies, each of which determines the manner in which the model acquires and applies its knowledge:

- **Supervised learning** trains the model on labelled images, allowing it to associate inputs with corresponding target image output. Throughout optimisation, the model learns to map input features to desired outputs. Iteratively, measuring and minimising the difference between predicted and actual outcomes.

- **Unsupervised learning** engages the model with unlabelled data where it must identify patterns or structured without explicit instructions.

[50] Justin D Weisz and others, 'Design Principles for Generative AI Applications', *Proceedings of the CHI Conference on Human Factors in Computing Systems* (ACM 2024) 2 <https://dl.acm.org/doi/10.1145/3613904.3642466> accessed 9 July 2025.

[51] Bandi, Adapa and Kuchi (n 26) 8.

[52] ibid 21.

- **Semi-supervised learning** blends both methods, using a small set of labelled data alongside a larger volume of unlabelled data. This approach can be particularly valuable when labelled data is difficult or costly to obtain.[53]

Irrespective of the approach adopted, human involvement remains essential. Developers are responsible for designing the training process, monitoring the model's progress, and validating its performance to ensure reliability and safety. [54] From a technical standpoint, the training process entails successive, gradual modifications to the model's weights and biases, as outlined in *Chapter 2.2*. These parameters, initially set to random values, are systematically updated as the model leverages feedback from the data. [55] The objective is to extract and internalise the underlying statistical relationships in the training examples, thereby enabling the model to generate images that are both novel and coherent in comparison to the training data.[56]

The selection of training algorithm significantly influences the training dynamics and the resulting generative capabilities. [57] Key variables in the training phase include the learning rate, batch size, and regularisation techniques. Fine-tuning these parameters is critical for achieving optimal performance and avoiding common pitfalls. [58]

Throughout the training process, the developer must make a number of choices, ranging from the selection of machine learning paradigms to the calibration of numerous hyperparameters. Each judgement determines how the model assimilates knowledge and how susceptible it becomes to pitfalls such as the pictures being too similar to the training data. As the training phase draws to a close, the focus naturally shifts to the subsequent task of evaluating and validating the model.

2.3.5 Evaluation and validation

The evaluation and validation phase is of critical important in the lifecycle of GENAI systems, since it ensures that models produce output that is not only technically robust but also reliable and aligned with intended use cases.[59] This phase directly influences the ultimate utility of the system. Overall, the process is multifaceted and often contains a range of evaluation metrics, depending on the use case.[60]

---

[53] European Union Intellectual Property Office (EUIPO) (n 21) 31.
[54] ibid.
[55] Kruse and others (n 31) 25.
[56] Bandi, Adapa and Kuchi (n 26) 8.
[57] European Union Intellectual Property Office (EUIPO) (n 21) 151.
[58] Bandi, Adapa and Kuchi (n 26) 9.
[59] European Union Intellectual Property Office (EUIPO) (n 21) 267; Negnevitsky (n 29) 6.
[60] Bandi, Adapa and Kuchi (n 26) 9.

2.3.6 Fine-tuning

The final stage of the development process, fine-tuning, is initiated when a foundation model is to be adapted for even more specific tasks, domains, or stylistic requirements. It is important to note that this step is not universally applicable to all development processes. Its implementation is selective and contingent on the specific demands of the intended application. Building on the broad and generalised capabilities established during the initial training, fine-tuning sharpens the model's performance toward the targeted use. This approach is particularly valuable for image generation, where the ability to specialise certain styles and/or objects often becomes essential.

Consider a photographer who has cultivated a unique personal style, portraying animals in their work. They now want to expand their portfolio by generating images that look as if they were taken with their camera in their own style, portraying imaginary animals. The model would require fine-tuning using the artist's existing portfolio, ensuring that the model becomes familiar with the patterns. Concurrently, the model is required to adhere to the inherent constraints of realistic photography, whilst simultaneously associating these with the representation of imaginary creatures. The realism of a depiction is contingent on the meticulous curation of a dataset, which establishes constraints on the depiction.

Utilising the training data that corresponds to the fine-tuning objectives, the objective is to once again adjust all or parts of the model's weights so that the model demonstrates a high level of proficiency in the new, more specific task, while simultaneously minimising the loss of the general knowledge acquired during the initial training. [61] Whilst fine-tuning generally necessitates a reduced amount of data in comparison to the construction of a model from the beginning, it does require higher-quality examples in order to ensure an accurate and authentic rendering of the desired style.[62]

[61] European Union Intellectual Property Office (EUIPO) (n 21) 148.
[62] ibid 125.

2.3.7 Summary

As previously discussed, the developer's choices at each stage of the development process are central to the legal analysis of possible authorship. The following key actions, available to the developer throughout the process, form the factual foundation for the legal evaluation:

Problem definition

- Determining the intended outcome or use case(s) for the system
- Defining software and hardware requirements

Data collection and preprocessing

- Selecting data sources and determining data origins
- Selection of training dataset composition and scope
- Deciding the dataset diversity and scope
- Developing and applying labelling strategy

Architecture selection

- Choosing the type and configuration of model architecture
- Specifying architectural parameters

Model training

- Selecting the training methodology
- Overseeing quality assurance and monitoring progress

Evaluation and validation

- Setting and applying standards for desired outcomes

Fine-tuning

- Specifying the target domain or style specification
- Curating supplementary or specialised datasets for further training

2.4 Generative process

As with the development process, this analysis traces available choices as these factors are crucial to the legal analysis. Although the user is said to have a plethora of options at their disposal, it is ultimately the developer's decisions that impose the primary limitation. This section will concentrate on this balance between user agency and developer control, as it is crucial to the possible authorship status, which is a question that will be discussed at a later point. As previously referenced, *This person does not exist* exemplifies a system in which the developer deliberately restricts user interaction to a single function, refreshing the page, while retaining all other generative controls. Generation begins automatically upon website loading, with users having only the option to restart the process by refreshing the page. Whilst it is true that alternative systems offer a greater range of interactive options, it is nevertheless crucial

to acknowledge that the implementation of these features is at the discretion of the developers, and they may choose to limit or remove them. This chapter clarifies the technical and practical boundaries of user and developer agency in the generative process, thereby establishing the necessary foundation for the subsequent legal analysis concerning the attribution of authorship for the generated images.

The generative process's starting point is the input. In the current era, prompts represent a predominant modality of input. The quality, relevance, and style of images are directly influenced by the specificity and clarity of a prompt. These instructions may encompass a wide spectrum, ranging from general artistic directives (*a painting of a dog*) to precise technical specifications detailing camera settings, lighting, or composition.[63]

It is important to note that the capacity to communicate with the system does not ensure that the image will invariably correspond to the user's instructions. The developer may impose limitations or remove this functionality through various mechanisms, thereby exerting an influence on the generative process. To mention a few; first, the developer's selection of training data may exclude certain characteristics. It is an established fact that models are only able to generate content based on patterns derived from their training. In the absence of specific associations, such as the link between for example the phrase *oil paint* and the imagery of paintings created in that medium, the model will be incapable of producing accurate representations. This limitation can be observed in public image generators, particularly with regard to unwanted content such as that portraying violence, hate, or adult material, which is deliberately minimised or excluded from training data.[64] Secondly, domain-specific fine-tuning has been demonstrated to increase the probability of the manifestation of certain characteristics. In the event of input that does not align with the system's intended purpose, the output image may be suboptimal or irrelevant. To illustrate this point, consider a system that has been meticulously calibrated to produce photorealistic images of individuals. In response to a request for *a realistic depiction of a person with a brown beard*, this system will generate images that is more aligned with the specified request than with a request for a cat in a *pixel art style*[65]. [66] Thirdly, it is possible for developers

[63] 'Prompt and Image Attribute Guide | Generative AI on Vertex AI' (*Google Cloud*) <https://cloud.google.com/vertex-ai/generative-ai/docs/image/img-gen-prompt-guide> accessed 11 July 2025.

[64] 'DALL·E 2' (*OpenAI*) <https://openai.com/index/dall-e-2/> accessed 11 July 2025.

[65] A digital art form where images are created and edited at the pixel level, often resulting in a retro, blocky aesthetic reminiscent of early video games.

[66] 'Controlled Customization | Generative AI on Vertex AI' (*Google Cloud*) <https://cloud.google.com/vertex-ai/generative-ai/docs/image/edit-controlled> accessed 11 July 2025.

to implement filters with the objective of preventing certain prompts from being processed or images from being displayed to the user.[67]

In addition to the prompt, developers have the capacity to empower users by enabling control over model predictability and variability. This can be achieved through the exposure of parameters such as the choice of seed[68], temperature[69], top-k[70], and top-p[71]. Alternatively, developers have the capacity to set these parameters themselves. This enables them to enforce consistent or variable images. In doing so, they may trade user agency for system reliability or safety.[72]

In summary, the user can have a significant impact on the generated images despite their absence during the development phase. However, it should be noted that the user can only access the capabilities through the user interface, the control of which is the responsibility of the developer. Consequently, the developer ultimately determines the manner in which the user interacts with the system. It is acknowledged that users are at liberty to provide input; however, it is imperative to recognise the impact of the temperature, top-k and top-p values on the system's response to the prompt. This precludes the impact that the development process has already had. Furthermore, the developer may employ tools such as filtering to further restrict the impact of the user's agency.

2.5 Summary

This examination of the development and generative process has yielded several conclusions, which will serve as the foundation for addressing the research questions posed in this thesis. The development process is comprised of numerous steps and decisions that must be made by the developer. These elements possess distinct characteristics and exert varying degrees of influence on the ultimate result. As will become clearer during the subsequent chapters, the balance of autonomy between the user and the developer is central to the argument. In essence, it is the developer who ultimately empowers

---

[67] 'DALL·E 2' (n 65).

[68] Refers to the initial random value used to determine the model's output generation process.

[69] Controls the level of randomness in token selection: lower values make the output more deterministic, while higher values increase creativity and variability.

[70] Restricts the model to sampling from the *k* most likely next tokens, thereby narrowing output diversity.

[71] Narrows the choice of next tokens to the smallest set whose combined probability mass exceeds *p*, ensuring a balance between coherence and diversity.

[72] Harald Semmelrock and others, 'Reproducibility in Machine-learning-based Research: Overview, Barriers, and Drivers' (2025) 46 AI Magazine 8 <https://onlinelibrary.wiley.com/doi/10.1002/aaai.70002> accessed 9 July 2025; Weisz and others (n 51) 3; 'Design Multimodal Prompts | Generative AI on Vertex AI' (*Google Cloud*) <https://cloud.google.com/vertex-ai/generative-ai/docs/multimodal/design-multimodal-prompts> accessed 11 July 2025.

the user to influence the images. However, the characteristics of GENAI complicate this relationship.

The system's opacity hinders the process of retroactive justification of the model's behaviour, thereby impeding the ability to discern the rationale behind a particular conclusion. To illustrate this problem in image generation, it is important to note that the exact reasons for the use of particular colours, the expression of a dog portrait, or the selection of a specific colour palette are not known. In certain instances, it is possible to formulate estimations, informed by factors such as the training data and the prompt, among other considerations. In the event of the training data having been constrained to a specific colour palette, it can be deduced that the rationale behind the model producing images in those colours is that it is based on the knowledge it has been given. In the alternative scenario, the input might comprise a specified colour palette. It can be deduced that the image is depicting the specific colour scheme as a consequence of the instruction.

To summarise, the chapter has examined the processes of both development and generation. These are of crucial importance in the examination of possible authorship. The subsequent chapter will examine the legal framework, focusing on the factors that are pertinent to the question of whether the developer can, or at least should, be granted status as an author. As previously mentioned in this chapter, particular emphasis will be placed on the available options.

3 Legal framework

3.1 Introduction

This chapter begins with *Chapter 3.1,* which identifies the fundamental sources of the EU copyright framework and summarises the key legal principles that influence the copyright landscape. In particular, the principles of proportionality and of a high level of protection. It then examines the requirements for copyright protection in Chapter 3.2, which are necessary to determine authorship. The requirements of originality and expression, to be specific. Furthermore, *Chapter 3.3* explores how collaborative contributions are treated, evaluating whether the user's chances for authorship limits the developer's claim to authorship. Further, Chapter 3.4 explores how past technological challenges put on the EU copyright framework. By studying the examples of photography and software, guidance can be found that applies to image generators. The final main section, *Chapter 3.5*, applies the aforementioned frameworks in systematic order to answer the first of the thesis's research questions: Under EU copyright law, to what extent can developers of generative AI systems be recognised as authors for images generated by their models, considering existing originality requirements and the concept of 'author's own intellectual creation' established in CJEU jurisprudence?

The right to one's intellectual property is protected as a human right by the *Charter of Fundamental Rights of the European Union* (the Charter) and confirmed by CJEU.[73] EU copyright legislation is fragmented and spread across several directives and pieces of jurisprudence. The most important legal text for the purposes of this thesis is the InfoSoc Directive, as it is one of the most impactful EU copyright directives and the closest thing to a basic rights codex within the framework.[74] Secondly, the case law from CJEU, as it is the ultimate interpreter of EU legislation.[75] With its crucial impact on the area, mostly from preliminary reference procedure, it can be argued that the precedence should even be viewed as law with its own authority.[76]

Furthermore, rights arise from national jurisdictions, and due to the limited harmonisation, variations exist within the EU. [77] For instance, moral rights

[73] Charter of Fundamental Rights of the European Union [2012] OJ C326/391 (the Charter), art 17(2); Case C-277/10 Martin Luksan v Petrus van der Let EU:C:2012:65.

[74] Maria Martin-Prat, 'Twenty Years After—the State of EU Copyright Law' in Hayleigh Bosher and Eleonora Rosati (eds), *Developments and directions in intellectual property law: twenty years of the IPKat* (First edition, Oxford University Press 2023) 131.

[75] Consolidated version of the Treaty on the Functioning of the European Union [2012] OJ C326/47, art 267.

[76] Ana Romalho, 'The Competence and Rationale of EU Copyright Harmonization' in Eleonora Rosati (ed), *The Routledge handbook of EU copyright law* (Routledge 2021) 12.

[77] Justine Pila and Paul Torremans, *European Intellectual Property Law* (2nd edition, Oxford University Press 2019) 223.

are not required under EU copyright law. However, they are present in most states due to obligations arising from Article 6bis of the Berne convention.

In addition to the sources previously mentioned, the EU copyright framework incorporates several principles. These standards can be derived from various sources, including the preambles of the directives or the Charter, for example. At other times, the CJEU relies on standards that it has developed itself or that are regarded as general principles of EU law.[78] two principles relevant to this thesis are proportionality and a high level of protection.

The principle of proportionality is rooted in the Charter, among other sources. As a fundamental right, copyright must be balanced as such against opposing rights. Examples of rights that often tend to compete with each other include the freedom to conduct business and freedom of expression and information.[79] The principle obliges legislators and courts to consider the interests of both right holders and users and the public, for example in questions of freedom of expression. This rights-based reasoning has become increasingly prevalent since the Charter became legally binding.[80] This balancing act is illustrated in the Deckmyn case, in which the CJEU clarified that exceptions such as the parody exception must be interpreted in light of the Charter. This requires courts to balance exclusive rights with artistic and social commentary.[81] This approach has been further developed and constitutionalised in more recent CJEU case law, most notably in Spiegel Online and Funke Medien.

In these cases, the Court held that all aspects of EU copyright law must comply with the Charter, including the scope of exclusive rights and the interpretation and application of statutory exceptions or limitations. This means that both national and EU authorities are obliged to ensure a fair balance between the rights of copyright holders and competing fundamental rights, such as freedom of expression and information. Thus, the decisions in Spiegel Online and Funke Medien reinforce the idea that copyright cannot be construed or enforced in isolation, but always in light of broader constitutional principles that protect fundamental rights.[82]

Moreover, the principle of high level of protection can be found in several locations within EU copyright law, most notably in recitals 4 and 9 of the InfoSoc Directive. As there are often competing values, as mentioned above,

[78] Rosati (n 11) 38.

[79] Pila and Torremans (n 74) 222; Tito Rendas, 'Fundamental Rights in EU Copyright Law' in Eleonora Rosati (ed), *The Routledge handbook of EU copyright law* (Routledge 2021) 20.

[80] Rendas (n 76) 19.

[81] Case C-201/13 Johan Deckmyn and Vrijheidsfonds VZW v Helena Vandersteen and Others EU:C:2014:2132.

[82] Case C-516/17 Spiegel Online GmbH v Volker Beck EU:C:2019:625; Case C-496/17 Funke Medien NRW GmbH v Bundesrepublik Deutschland EU:C:2019:623.

one might ask which object should be given high protection. In case law, the object to be protected varies. The CJEU has applied the principle when examining the scope of economic rights, as well as enforcement and remedies. It has been used to support arguments for expansive interpretations of copyright rights, as well as to protect rightsholders and authors specifically.[83]

3.2 Construction of copyright

The following section unpacks the core legal test for copyright eligibility under EU law, examining the requirements of originality and expression. It establishes the boundaries of what constitutes a work and who can claim authorship, providing a legal basis for subsequent analyses of the relationship between generated images and the developer of the system that created them.

Despite never being explicitly defined in the legislation, the work and the author are the two central factors in copyright law. Instead, well-established principles must be relied upon to decipher the concepts. These two concepts are autonomous in EU law and must be applied with a uniform interpretation.[84] Although it might seem counterintuitive, the concept of the author is interdependent on the concept of the work.[85] The work is the focal point of protection when it is an original expression. This interpretation represents the two thresholds that the CJEU has upheld, making it clear that works must not be subject to any other requirements. These standards present the limitations of what can be protected under EU copyright law, a topic that will be expanded upon in the following chapter.[86]

3.2.1 Originality

While opinions differ on whether the criterion of is easy to uphold, it is clear that the same threshold must be applied to all subject matter.[87] The concept has been fleshed out in many cases in the CJEU, where it has been the direct or indirect subject. The term should be understood as meaning that the subject matter has been derived or originated from the author.[88] Its purpose is to ensure a sufficient link exists between the work and the author.[89] However, creative elements must also be added, even though creativity is not a requirement

[83] Rosati (n 11) 41.

[84] Case C-227/23 Kwantum Nederland BV and Kwantum België BV v Vitra Collections AG EU:C:2024:844, para 56.

[85] Aurelija Lukoševičienė, *Authors and Creative Users: Addressing the Conceptual Challenges of Digital Creativity to EU Copyright Law from the Perspective of the Author* (Lund University (Media-Tryck) 2021) 86.

[86] ibid 223.

[87] Eleonora Rosati, 'Why Originality in Copyright Is Not and Should Not Be a Meaningless Requirement' (2018) 13 Journal of Intellectual Property Law & Practice 597, 1; Stef van Gompel, 'Creativity, Autonomy and Personal Touch: A Critical Appraisal of the CJEU's Originality Test for Copyright' in Mireille MM van Eechoud (ed), *The Work of Authorship* (Amsterdam University Press 2015) 95; Lukoševičienė (n 82) 226.van Gompel 95.

[88] Case C-145/10 Eva-Maria Painer v Standard VerlagsGmbH and Others EU:C:2011:798, para 87.

[89] Pila and Torremans (n 74) 74.

per se. The term *original* must not be confused with its meaning in other fields of science; in the arts, for example, there is a significant discrepancy. In law, the requirement is measured solely on the subject matter in question, whereas in most other disciplines it is treated as a relative or comparative assessment.[90] In the legal sense, original does not mean new or never seen before. There is no threshold for the subject matter to have a certain aesthetic characteristic or to be culturally significant.[91] Instead, the primary concern lies in the link between the author and the work. This becomes especially clear when it comes to the protection of creations such as football fixture lists, algorithms, and the format of data files in computer programmes.[92]

Over the years, the CJEU has described the link between work and its author in various ways. The first landmark case was Infopaq. The Court ruled that the work must be the author's own intellectual creation.[93] This is a criterion that now encapsulates the EU copyright standard.[94] Subsequently, the question was revisited in Painer, where the Court emphasised that an intellectual creation is an author's own if it reflects their personality. By making *free and creative choices*, the author would *express their creative abilities* and thus *reflect their personality in their work.* As previously mentioned, this phenomenon has been described in various ways. Most notably, is has been described as leaving a *personal imprint* or putting their *personal stamp* or *touch* on the work. [95] This should not be confused with having an easily detectable signature that someone familiar with the artist might recognise.

The CJEU has set out a two-step process for determining whether a work is the author's own intellectual creation.[96] Firstly, determining whether the medium allows sufficient room for creative choices to be made. While few subject matters do not fulfil this requisite, it has been seen in individual words and pixels of colour. This freedom can also be limited by constraints, such as technical considerations. Free and creative choices are not made when the actions are dictated by technical or functional considerations.[97] However, this

[90] van Gompel (n 84) 137.

[91] Case C-683/17 Cofemel – Sociedade de Vestuário SA v G-Star Raw CV EU:C:2019:721.

[92] Joined Cases C-403/08 and C-429/08 Football Dataco Ltd and Others v Yahoo! UK Ltd and Others EU:C:2012:115, paras 29-45; Case C-406/10 SAS Institute Inc. v World Programming Ltd EU:C:2012:259, para 45.

[93] Case C-5/08 Infopaq International A/S v Danske Dagblades Forening EU:C:2009:465, para 38; Case C-145/10 Eva-Maria Painer v Standard VerlagsGmbH and Others EU:C:2011:798, para 87.

[94] Marianne Levin, 'The Cofemel Revolution - Originality, Equality and Neutrality' in Eleonora Rosati (ed), *The Routledge handbook of EU copyright law* (Routledge 2021) 86.

[95] Case C-145/10 Eva-Maria Painer v Standard VerlagsGmbH and Others EU:C:2011:798, paras 88-92. Joined Cases C-403/08 and C-429/08 Football Dataco Ltd and Others v Yahoo! UK Ltd and Others EU:C:2012:115, para 38.

[96] Case C-145/10 Eva-Maria Painer v Standard VerlagsGmbH and Others EU:C:2011:798.

[97] Case C-833/18 SI, Brompton Bicycle Ltd v Chedech/Get2Get EU:C:2020:461.

does not completely exclude the entire work from protection. Instead, the elements resulting from these choices should be disregarded from examination.[98] Constraints can also be imposed by other factors, such as the choice of genre. As constructed categories of creative pieces, some choices are interpreted as not being the author's own expression. Therefore, elements that are associated with the chosen style are must be disregarded.[99]

Secondly, the question of whether or not the room has been used is analysed.[100] This must not be confused with the idea that the author must demonstrate a certain level of effort or skill in the project.[101] Instead, the assessment is qualitative rather than quantitative. The assessment focuses on the author's process and the impact on the result, rather than the additions made in terms of effort, skill or other qualities.[102]

In the specific case of Painer, the Court was stated that even a portrait photographer has the possibility of making the necessary choices to reflect the author's personality. Whether the action was taken in this case was left for the national court to determine as part of the second stage of the analysis. However, the CJEU stated that free and creative choices can be made in many ways and at various stages of production. In the case of the portrait photographer, the creative production process was divided into three phases: preparation, taking the photograph, and selection of snapshot. In the first phase, the author can choose the background, subject's pose and lightning. In the second, the framing, angle of view, and create atmosphere. Lastly, the choice of developing technique and possible usage of computer software. The Court made it clear that with the three phases, the author indeed has freedom to exercise his creative abilities in a way that can result in an original work. the case does not mention any hierarchy between the stages, doctrine suggests that choices made later in the production process, especially the final decision, seem to be important, at least in large collaborative projects.[103] Having the final say in productions can often be assumed to leave a personal imprint, particularly for creative leads.[104]

3.2.2 Expression

Under EU law, the criterion of expression determines which subject matters are eligible for copyright protection. This requirement can be described as the link between an author's creative act and the manifestation that is left behind.

---

[98] Case C-393/09 Bezpečnostní softwarová asociace – Svaz softwarové ochrany v Ministerstvo kultury EU:C:2010:816.

[99] van Gompel (n 84) 112.

[100] Case C-145/10 Eva-Maria Painer v Standard VerlagsGmbH and Others EU:C:2011:798.

[101] Joined Cases C-403/08 and C-429/08 Football Dataco Ltd and Others v Yahoo! UK Ltd and Others EU:C:2012:115.

[102] Lukoševičienė (n 82) 241.

[103] Pila and Torremans (n 74) 255.

[104] van Gompel (n 84) 138.

The CJEU has confirmed that the work must be expressed in a form that can be identified with sufficient precision and objectivity.[105]

Article 2(1) of the Berne Convention stipulates that all literary and artistic works, "whatever may be the mode or form of its expression", are eligible for protection.[106] However, protection is limited to the specific, concrete expression of an idea, not to the underlying idea itself. The Berne Convention explicitly states that only expressions can be protected, not ideas, procedures, methods of operation or mathematical concepts. Although this principle has not been explicitly mentioned in EU legislation, the CJEU has nevertheless applied it.[107] This critical distinction is crucial for determining the scope of protection and preventing monopolisation of underlying ideas.

An expression does not require physical manifestation, but it must exist in a form that enables the work to be "identifiable with sufficient precision and objectivity". The CJEU used this wording in the Levola Hengelo case when examining whether taste could be protectable subject matter.[108] The Court stated, that at the time, was not possible to objectively describe a taste and thus qualify it as an expression. By contrast, a song is a subject matter that is not physical but can be identified with precision due to objective parameters such as melody, rhythm and lyrics. This issue was also discussed in the Cofemel case, in which the focus of the assessment was the aesthetic effect.[109] As with taste, aesthetics were regarded as a subjective element that can be perceived differently by different people. Perception may differ between individuals and even within an individual, and this may be influenced by circumstances.[110]

The need to objectively identify the work seems to stem from the requirement for other parties to understand what the work entails and its limitations.[111] This has practical implications for avoiding and establishing infringement, and other protective measures.[112] Protection extends beyond any physical manifestation to the immaterial dimension, encompassing what is subjected beyond the mere idea.[113] It is important to distinguish between subject matters that have touchpoints but should be examined as separate entities. For example, a video game can consist of many copyrighted assets, such as music

[105] Case C-310/17 Levola Hengelo BV v Smilde Foods BV EU:C:2018:899.
[106] Berne Convention art 2(1).
[107] Lukoševičienė (n 82) 248.
[108] Case C-310/17 Levola Hengelo BV v Smilde Foods BV EU:C:2018:899, paras 35-43.
[109] Case C-683/17 Cofemel – Sociedade de Vestuário SA v G-Star Raw CV EU:C:2019:721, para 53, 48.
[110] Levin (n 91) 88.
[111] Lukoševičienė (n 82) 253.
[112] ibid 254.
[113] ibid 223.

tracks, character designs and dialogue scripts, not to mention cinematic sequences.[114]

3.3 Joint authorship

This section clarifies how EU copyright law treats works created by multiple contributors, with a specific focus on joint authorship. By examining the circumstances in which collaborative contributions may satisfy the originality requirement, the relationship between the user and the developer is illuminated.

The core concept of joint authorship is not harmonised within the EU. However, it is not an unfamiliar concept, as it is mentioned in *Directive 93/98/EEC harmonising the term of protection of copyright and certain related rights* (Term Directive). [115] The text explicitly reveals that the lifetime of a work created through joint authorship is defined. From this acknowledgement, and others like it, one can conclude that the concept is not precluded within the Union. This is supported by the opinion of the Advocate General in the pending preliminary ruling case Chabrol C-182/24 although validity is not the focus. Consequently, it is left to the Member States to regulate collaborations. [116] Despite the lack of harmonisation, some argue that general principles based on interpretations of the copyright framework should lead to joint authorship. One such scenario is when the contributions of each author to a work cannot be separated, yet each author's input would satisfy the originality criteria independently. [117]

Naturally, there are also scenarios where multiple people are involved, but not all of them are authors. In many cases, it can be difficult to discern the level of an individual's involvement in the work. A clear hierarchy can make it easier to identify creative leaders, but a lack of hierarchy should not necessarily preclude originality. [118] Works with easily separable parts also make this examination easier. For example, consider a song where the guitarist wrote their part, the drummer wrote theirs, and the singer wrote the lyrics. Compared to a scenario where all three people worked on all three parts of the song, the first alternative makes for an easier evaluation, since each person's involvement is clearly defined. In the latter type of collaboration, some argue that making decisions later in the process is advantageous for fulfilling the originality criteria. This is because their personal touch would not have been lost, as the fewer choices made after their involvement, the lower the

[114] Case C-355/12 Nintendo Co. Ltd v PC Box Srl and 9Net Srl EU:C:2014:25.

[115] Council Directive 93/98/EEC of 29 October 1993 harmonising the term of protection of copyright and certain related rights [1993] OJ L290/9 (Term Directive), art 6 and recital 17.

[116] Giulia Priora, 'Copyright Law and the Promotion of Scientific Networks: Some Reflections on the Rules on Co-Authorship in the EU' (2019) 9 Queen Mary Journal of Intellectual Property 217, 222.

[117] Lukoševičienė (n 82) 348.

[118] van Gompel (n 84) 131.

risk. From a practical perspective, the method must offer a relative stability to preserve the link between the artist and their work. Depending on the medium, artistic input may be erased or developed in such a way that it becomes indiscernible.[119]

The type of input can vary greatly depending on the medium. This can range from a more direct, hands-on approach to artistic guidance. For the sake of argument, some even consider that a painter's community can have a sufficient impact, given that their influence can affect the painter's technique, resulting in specific art pieces.[120] However, this should be taken lightly and viewed as an intellectual thought experiment on the limits of joint authorship. The point is that co-authorship can take many forms, depending on factors such as the nature and closeness of the collaboration. Crucially, the creative choices must be evident in the final product. The stronger the collaboration, the greater the likelihood that the work will reflect a joint personal imprint.[121]

In summary, this section demonstrates that, while joint authorship is not harmonised at EU level, it remains legally possible when individual contributions are inseparable yet independently original. The allocation of authorship between developers and users of GENAI therefore hinges on the nature and extent of each party's creative choices and requires careful, case-by-case analysis under EU copyright law.

3.4 Historical challenges

This section provides a contextual framework for understanding the challenges posed by GENAI by examining the historical responses of EU copyright law to technological shifts, with a particular focus on the introduction of photography and software. By analysing the transition from initial scepticism to eventual recognition of authorship in these media, the chapter elucidates analogies and lessons for addressing GENAI, demonstrating that established legal principles can adapt flexibly to copyright questions raised by emerging technologies.

The EU copyright system has a long history, and the legislation should, as far as possible, allow for a coherent interpretation when new technology emerges. In order to ensure the correct application of the framework in accordance with its intended purpose, it is imperative that innovation is supported and that the framework does not act as an impediment.[122] It is evident that technological advancements have had a significant impact on the copyright system. This has resulted in the need for reassessments of core principles

[119] Pila and Torremans (n 74) 72.
[120] ibid.
[121] van Gompel (n 84) 133.
[122] Pila and Torremans (n 74) 3.

and reinterpretations to establish the necessary flexibility. In other instances, the process may be characterised as a mere crystallisation and clarification.[123]

The current view of photography as a medium has already been discussed in this chapter, at least regarding portraits. However, when the question of eligibility first arose, it was contested due to the contrast with what had previously been protected. [124] One of the barriers was that, from an outsider's perspective, the art form only required the push of a button, and thus photographers would rather be seen as technicians than artists. Consequently, a debate arose about the source of the work itself. The entity responsible is either the camera or the human operating it.[125] As previously mentioned in this chapter, photographs are now explicitly stated as eligible subject matter for protection in the Term Directive. Additionally, the CJEU set out the precedent for carrying out the originality assessment in Painer, as mentioned previously. It was emphasised that photographers have considerable scope for free and creative choices. The camera is merely a tool through which individuals can express their creativity. A parallel can be drawn to the developer, who like the photographer, has an indirect connection to their images. This point will be further developed in the next main section of the thesis.

Software is a more recent example of a technology that has disrupted the copyright system. Notably, two factors set it apart as a subject matter. Firstly, its technical nature. Software in the form of source code is not intended for human enjoyment in itself. Instead, the purpose it to provide it to a machine for translation and have it perform according to instructions. Although understandable, the code has not an intuitively creative purpose. Secondly, even though the software can objectively be viewed as a text, its purpose is functional. This intent differentiates it from other mediums, and it could be compared to instructions, which are not eligible for protection.[126] This functional purpose blurs line between idea and protectable implementation, challenging the idea/expression dichotomy. Additionally, the author's free and creative choices of the author are often cases obscured. Compared to works such as photographs, it is generally difficult to discern the personal touch from the examining the work.[127]

Today, software is recognised as copyright eligible subject matter under the InfoSoc Directive and is also being explicitly mentioned in *Directive*

[123] Lukoševičienė (n 82) 158.
[124] ibid 147.
[125] ibid 148.
[126] ibid 158.
[127] ibid 168.

2009/24/EC on the legal protection of computer programs (Software Directive) under the term computer programs.[128] These directives provide similar rights and threshold for protection. However, if the software falls outside the scope of the Software Directive, it is unlikely included in the InfoSoc Directive due to their overlapping scopes. [129] The definition of *computer program* is held broad, encapsulating all forms (both source code and object code) as well as preparatory design material.[130] However, case law has made it clear that the programming languages are excluded as well as graphical interfaces. The scope is limited to cases where reproduction of the part would result in the reproduction of the programme itself, meaning that it would perform the same tasks.[131]

To summarise, the historical responses to photography and software reveal that EU copyright law is capable of accommodating new technologies by focusing on the underlying human choices and expressions, rather than the medium itself. These analogies demonstrate that, while technological advances may initially disrupt established doctrines, EU law has demonstrated the capacity to reinterpret core principles to encompass novel forms of creativity. Accordingly, the lessons derived from the analysis of photography and software provide a rational foundation for addressing the copyright challenges posed by GENAI. This suggests that established legal standards can be adaptably applied as new questions of authorship and originality emerge.

3.5 The developer scenario

This final main section of *Chapter 3* is dedicated to the primary research question: Under EU copyright law, to what extent can developers of GENAI systems be recognised as authors for images generated by their models, considering existing originality requirements and the concept of 'author's own intellectual creation' established in CJEU jurisprudence? Drawing upon the technical foundations established in *Chapter 2* and the legal framework previously examined in *Chapter 3*, this analysis systematically evaluates the authorship criteria, identifying the strengths and potential obstacles of developers in claiming copyright protection. The analysis emphasises actions that can enhance the developer's eligibility for authorship, while also acknowledging points of friction within the current legal framework.

3.5.1 Originality

Whilst the originality criterion is widely regarded as being uncomplicated in terms of its satisfaction, it presents particular challenges for developers of

[128] Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs [2009] OJ L111/16 (Software Directive), art 1.

[129] Pila and Torremans (n 74) 259.

[130] Software Directive art 1.

[131] Case C-393/09 Bezpečnostní softwarová asociace – Svaz softwarové ochrany v Ministerstvo kultury EU:C:2010:816, paras 48-49.

GENAI systems. This fundamental requirement establishes a link between the work and the author by recognising the free and creative choices that express the author's personality. In essence, this discussion will focus on whether the development process allow the developer to make free and creative choices in regards to the images that are generated by their systems. The analysis will first identify the choices that the developer makes and secondly consider whether these choices in fact does leave a reflection of the developer's personality in the images.

3.5.1.1 Scope of free and creative choices

The initial hurdle that developers encounter is analogous to the historical prejudice that has been faced by software developers and photographers. Specifically, there is a perception that their work is merely technical in nature and lacks creative expression. In a manner analogous to the operation of a camera by a photographer through the simple pressing of buttons, the developer's actions may appear to be the mere inscription of code. However, this superficial perspective fails to acknowledge the intricate nature of these processes, which necessitate a multitude of decisions and considerations. The Painer judgment is notable for its recognition of the cumulative effect of creative choices, acknowledging that even actions which appear to be mechanical in nature can involve significant creative discretion when viewed in their proper context. Nevertheless, it should be noted that the process is subject to a number of technical considerations of varying degrees of complexity. The copyright framework does not demand a binary categorisation of actions as either creative or technical. As emphasised in case law, choices are free and creative when they are not solely dictated by technical considerations. This nuanced approach is of crucial importance when evaluating the developer's process, which combines creative vision with technical implementation. Decisions can be categorised into three distinct types: those that are purely creative, those that are creative but are influenced by technical factors, and those that are purely technical. This categorisation provides a framework for analysis, while acknowledging that the boundaries are not always clear-cut.

Creative	
Identification of desired outcome for the model	Requires judgement, alignment with intentional output or at least use-case. This shapes the system's purpose and informs subsequent chocies. For small task-specific models that are made to produce specific imagery, here is where the planning starts to make the vision come to life.
Selection and exclusion of training data	Requires judgement that aligns with the intended outcome. Deciding which artists, genres and objects that should be portrait

	in the sets. While a foundation model requires considerations to make sure that they get exposed to a wide variety of data, it is still not something that depends on functionality per se.
Specification of target style or domain	Developer imposes stylistic boundaries and thematic focus, which infuses the system's image output with a particular *vision*. Not simply responding to functional needs, but to an expressive ambition. This choice can be cemented so that the model won't produce anything outside of this area of expertise, for example if the developer chooses to filter the images.
User interface design	Decisions about UI determine the user's degree of agency and shape how the system is framed as a creative tool. This reflects the developer's unique philosophy about creativity and user involvement

Creative but not without technical considerations

Choices of data sources	While sources must meet some legal and data-format requirements, significant aesthetic/curatorial judgement is exercised in choice of datasets (e.g. institutional archives vs. popular media)
Setting the desired diversity and scope of the data set	While having technical considerations such as making sure to hit certain targets, there is a large discretion for creativity regarding breadth, niche etc. Closely related to the choices in the problem definition
Labelling strategy	Approach must suit ML requirements, but allows creative prioritisation of specific content, tags, or semantic relationships
Selection of architecture	Different architectures (GANs, VAEs, DMs) confer different expressive affordances; technical trade-offs exist, but

	the direction is chosen partly for artistic impact
Evaluation and validation	Metrics are selected to correspond with project objectives, balancing technical optimisation with qualitative goals (e.g., "creativity," "coherence")
Filtering implementation	Filters may be guided by technical/safety standards, but the particular implementation, what is restricted, and how, is shaped by developer priorities (having the model produce only in a certain style, specific objects)

Purely based on technical considerations

Specifications of hardware and software requirements	Dictated by computational needs, system compatibility, and efficiency. These choices lack room for expressive or artistic intention
Macro level specifications	Although impacting outcomes, these decisions are driven by engineering constraints and best practices rather than creative aims
Data preprocessing and cleaning	Removing errors/outliers is essential for system function and reliability, not for shaping expressive character

It is evident from this presentation that the developer has the capacity to exercise numerous choices in accordance with their creative vision. However, the characteristics of GENAI complicate the matter, as EU copyright law also requires a causal relationship between the author's intellectual effort and the work. This requirement is a particular challenge, insofar as the effort has a direct link with the system, but only an indirect connection with the images produced. In contrast to photographers, whose choice during capture results directly in identifiable elements within the final work, developers face the potential lack of causation. This challenge will be discussed in the subsequent section.

3.5.1.2 Causal connection between choices and the generated images

As demonstrated in Chapter 2.4, an analysis of the generative process reveals that the developer's influence on the images is subject to variation. The strength of the causal connection is contingent on the type of model produced and the user agency. In the context of systems engineered to generate a wide array of imagery with significant user control, it becomes challenging to assert that all generated images accurately reflect the developer's creative vision. Nevertheless, in the context of highly specialised systems with limited image variability, the connection between developer choices and final expression becomes more apparent. The principle is illustrated by the example of the photographer from *Chapter 2.3.6* who develops a system specifically to generate images of imaginary creatures in their distinctive style. When the developer fine-tunes the system according to their preference, selects representative pictures, and implement filters that does not allow for anything else to be depicted than images of imaginary creatures in the photographer's characteristic style, they can achieve control comparable to traditional photographic practice. With this control, one can presume that the images do reflect the developer's personality.

A critical question emerges regarding whether developer choices must be visually identifiable in the final work. The choices mentioned by the Court in Painer can generally be identified when viewing the photograph, yet this identifiability is not explicitly stated by the case law. The CJEU has constructed exceptions to originality narrowly, allowing choices to be partially dictated by technical considerations while excluding only those entirely determined by such factors. This approach suggests that the capacity for creative choice, rather than the visibility of specific choices in the final work, constitutes the relevant criterion. The correct interpretation seems to be that free and creative choices need not be identifiable in the work itself. Architectural choices that significantly impact generation but remain invisible to viewers should still qualify as free and creative when not entirely guided by technical considerations. This interpretation aligns with the framework's rejection of aesthetic criteria and its inclusive stance toward various forms of creative expression.

3.5.1.3 Reflections on originality

To summarise, the development process does entail a significant amount of free and creative choices; however, these may not be present in the images generated by the system. It is evident that the most significant choices pertain to the definition of the problem, the specification of the target style or domain, and the degree of user interaction. While the choices must not be visually identifiable, it is essential that they exhibit a causal connection to the developer's choices in order to reflect their personality. The strength of the connection is subject to variation depending on the user agency involved. In cases where the user has a higher degree of autonomy, there is a greater probability

of the developer's impact being compromised. Exemplification of this phenomenon is provided by systems such as DALL-E, Midjourney and Stable Diffusion. Nevertheless, the developer may limit the user's autonomy through mechanisms such as the filtration of generated images and the restriction of training data.

3.5.2 Expression

This section is dedicated to the expression criterion, which stipulates that works must be manifested in a form that can be identified with sufficient precision and objectivity. Although images in themselves do not present any inherent obstacles to meeting this requirement, the question remains as to whether AI-generated images can be considered expressions of the developer's creative intent rather than mere technical artifacts produced through statistical processes.

3.5.2.1 Temporal disconnection and causal mediation

A fundamental challenge arises from the causal disconnection between the developer's creative actions and the manifestation of the final image. Unlike traditional creative processes where artists maintain direct control over expression, generation depends on pattern recognition and statistical reproduction by the system, occurring without direct developer intervention. The developer prepares and designs the system during the development process, but the actual generation of images requires user initiation and system execution.

This temporal separation creates a multi-layered expression problem analogous to, yet distinct from, the photography precedent established in Painer. Where photographers exercise creative choices across three temporal phases (preparation, capture, and selection) developers face a more complex scenario where their preparatory choices must be sufficiently embedded within system architecture to survive both temporal disconnection and user intervention. The critical question becomes whether the developer's creative vision can be imbedded within the system's structure sufficiently so that the images it produces will carry it.

3.5.2.2 System specificity and expressive strength

The strength of expression varies significantly with system specialisation. Highly constrained systems, such as the previously mentioned photographer's imaginary creature generator, demonstrate stronger expressive connections than general-purpose models like DALL-E or Midjourney. Where developer choices create narrow creative parameters that consistently channel generation toward specific aesthetic outcomes, the temporal disconnection becomes less problematic for expression analysis.

This graduated approach aligns with EU copyright law's flexible accommodation of new technologies demonstrated in photography and software precedents. The framework's capacity to recognise expression across diverse media suggests that statistical generation processes can satisfy expression requirements when developer choices meaningfully constrain and direct generative possibilities, even without direct manipulation during the generation phase.

3.5.2.3 Reflections on expression

The expression criterion is met when a developer's creative choices during system design meaningfully shape the range and character of AI-generated images. EU law accommodates such indirect authorship so long as the developer's vision is embedded in algorithmic architecture, training data and parameter constraints in a way that affects the generated images. The more narrowly a system is crafted to channel images toward predetermined aesthetic outcomes, the stronger the expressive link to its developer. Conversely, general-purpose models, giving the user greater agency, exhibit weaker connections. In all cases, fine-tuning and filtering that enforce stylistic or thematic boundaries bolster the argument that generated images can constitute the developer's original expression.

3.5.3 Joint authorship

This final section of the chapter will discuss whether the user's claim to authorship limits the developer's possibilities. Based on the previous examination, user co-authorship does not automatically preclude developer authorship claims. However, it does raise questions about how creative contributions should be allocated when multiple parties influence the final images. Should joint authorship be relevant, both the developer and the user must satisfy the requirements for recognition as authors separately. The user's agency may dilute or erase the developer's personal imprint, particularly when interacting with systems capable of producing diverse imagery based on detailed user specifications. However, the extent of this dilution depends entirely on the developer's choices during system design. When developers limit user agency through filtering, constrained generation parameters or specialised training, their creative imprint may remain dominant regardless of user involvement.

The degree of user control varies dramatically across different AI systems, creating a spectrum of authorship possibilities. At one extreme, users can exercise detailed control over the generation process through prompt engineering, parameter adjustment and iterative refinement. When systems respond precisely to user instructions and allow fine-grained creative control, the situation becomes comparable to using graphic design software such as Photoshop, where the tool enables but does not create artistic expression. For example, imagine a user instructing a system to produce a grid with specific colour codes for each square. This level of control is similar to direct digital art creation, where the user's creative choices determine virtually every aspect

of the final image. In such cases, the developer's role may be limited to providing a sophisticated tool rather than contributing to the creative process itself. At the other extreme, systems that operate with minimal user input, predetermined generation parameters, and constrained dipiction possibilities maintain stronger connections to developer creativity. The developer's choices regarding training data, architectural design, and output filtering continue to influence every aspect of generation, regardless of limited user involvement.

The joint authorship analysis reveals that developer authorship claims are strongest when the system design restricts user freedom while directing generation towards specific creative outcomes that reflect the developer's vision. Therefore, developers seeking copyright protection should consider implementing design choices that maintain creative control while providing users with tools for exploration, rather than granting them comprehensive creative authority. However, a user's claim to authorship does not preclude the developer's claim.

4 The theoretical elements of authorship

4.1 Background

The following chapter provides the normative and theoretical basis for the thesis, evaluating whether the developer should be recognised as the author of the generated images. It examines classical justification theories (utilitarian, labour and personality theory) which offer distinct rationales for granting copyright. After outlining each theory, *The developer scenario*-sections applies the respective theory to the developer's situation. The aim is to assess whether the developer's role aligns with the foundational principles of each theory. I n doing so, this section directly addresses the second research question of the thesis, offering a rigorous theoretical framework with which to evaluate if extending copyright to developers is normatively warranted within the EU context.

Justification theories have been developed to explain the rationale behind granting individuals limited monopolies over their work. They are multifaceted tools that can be used to interpret and critique valid and historical law.[132] Copyright is not a natural occurrence in the free market, but rather a constructed system designed to regulate the competition. This system balances the tension between the authors' claims and the public's interest in accessing creations, and it requires justification.[133] As a legal concept, copyright aims to protect the products of specific creative acts. It is connected to the achievement of creation rather than to a physical piece. However, expression is undeniably required.[134]

This overview does not aim to provide an exhaustive understanding of the countless interpretations and sub-categories. Instead, it presents the most prevalent schools of thought in EU copyright law and explores the extent to which they support authorship for developers.[135]

4.2 Utilitarian theory

4.2.1 Introduction

The theory is based on the concept that an action is measured by the total net utility that it brings to the group of interest, which is calculated by summing the pleasure and subtracting the pain. However, pleasure and pain can be interpreted differently depending on the circumstances.[136] An exclusively utilitarian perspective is based on the idea that the only reason to grant exclusive

[132] Lukoševičienė (n 82) 185.
[133] Pila and Torremans (n 74) 6; Ang (n 5) 130.
[134] Ole-Andreas Rognstad, *Property Aspects of Intellectual Property* (Cambridge University Press 2018) 52.
[135] ibid 90.
[136] ibid 13.

rights to the creator is for the greater good. Assuming it is not a net positive legal action, the creator has no more justifiable claim to their creation than a third party. [137] To justify authorship for the developer, the net utility must be maximised. However, a fundamental issue with the theory is that it is difficult to apply. In most cases, pleasure and pain, and thus net utility, are notoriously hard to quantify.[138]

In the case of copyright, it is the general public who bear the most prevalent cost of the author's monopoly since they are deprived of the rights reserved for the author. [139] However, these exclusive rights are considered justified because they are expected to lead to maximise net utility gain by providing an incentive to create works. Furthermore, copyright in the EU contributes to many of the Union's other interests, such as respecting and promoting cultural diversity and freedom of expression.[140] Despite extensive research, there is limited evidence to support the effectiveness of exclusive rights as an incentive. However, no other model has been shown to produce better results. It is important to note, however, that significant differences can occur depending on the sector. In some markets, the initial seller benefits to such an extent that long-term exclusive rights become overprotective. One example of this is the computer industry, where products have a short shelf life and competitors will lag behind, leaving a free market window for the initial creator.[141]

Despite the difficulties of quantification, there are approaches that can be used to quantify the cost-benefit equation. The law and economics movement has drawn on the public good problem over the last 70 years, with a particular increase in recent decades. This method uses economic theories to equate the economic consequences, aiming to create an optimal incentive to produce.[142] However, the utilitarian approach can encompass values other than economic ones. For example, the effect-based approach focuses on social efficiency, encompassing values such as education and freedom of expression.[143]

Other than providing incentives for creation and supporting the EU's other objectives, the concept of copyright is intended to prevent market failure. The inevitability of market failure is based on several behaviours that are presumed to be observed on a free, unrestricted, market. As authors require incentive to create, a lack thereof incentive and/or financial compensation for the creation leads to market failure. In order to maximise net utility, the work

[137] Pila and Torremans (n 74) 20.
[138] Rognstad (n 130) 16.
[139] Lukoševičienė (n 82) 190.
[140] Rendas (n 76) 20.
[141] Pila and Torremans (n 74) 20.
[142] Lukoševičienė (n 82) 187; Rognstad (n 130) 20.
[143] Lukoševičienė (n 82) 188–189.

should be subjected to the treatment that avoids these scenarios.[144] Three interconnected market failures that arise once a creative work is disclosed will now be presented.

The first is the problem of economies of scale whereby the producers in a competitive market face high, fixed costs when producing creative works. On a free market, one would expect to see near-zero margins, meaning that the originator will not be able to recoup the costs without protection.[145] Secondly, there is the public good problem whereby in contrast producers do not have the same costs. Since creative works are non-rivalrous and effectively non-excludable, subsequent entrants can reproduce the content at a negligible cost. This is the problem that has been most prominently recognised in the market of creative works.[146] Thirdly, the free rider problem occurs when an individual copies the creative efforts of others. This leads to a decreased incentive amongst creators, resulting to welfare loss.[147]

Traces of utilitarian argumentation can be found in the InfoSoc Directive as well as *Directive 2019/790 on copyright and related rights in the Digital Single Market* (DSM Directive), as well as un case law.[148] The presumption is that copyright incentivises the creation and exploitation of creative content for the benefit of both authors and the public, in addition to culture at large.[149] This implies a utilitarian view, as the justification of rights is based on serving not just the artist but those of the general public.[150] The use of exclusive rights as an incentive and means of financial protection of creators is also evident.[151] The same incentive-based reasoning can also be found in the DSM Directive and CJEU case law.[152]

4.2.2 The developer scenario

As established in the preceding section, utilitarianism quantifies the legitimacy of rules in accordance with the net social welfare that they yield. The concept of social welfare is inherently multifaceted and is subject to varying considerations depending on the specific context. Within the European Union, the primary concerns that inform the framework of author eligibility and copyright pertain to the promotion of cultural diversity, the safeguarding of freedom of expression, and the facilitation of a functioning internal market.

---

[144] ibid 190.
[145] Rognstad (n 130) 17.
[146] ibid 18.
[147] ibid 20.
[148] 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 [2019] OJ L130/92 (DSM Directive).
[149] InfoSoc Directive recitals 2 and 9.
[150] Rognstad (n 130) 89.
[151] InfoSoc Directive recital 10.
[152] DSM Directive recital 2.

The developer's claim to authorship would undeniably have a significant impact on multiple variables. In order to facilitate a meaningful analysis, the ensuing discussion will concentrate on the incentivisation of production of image-GENAI systems and GENAI-produced images, the impact on market structure, the effect on cultural diversity and the impact on free expression. Furthermore, the analysis must consider the principles of EU law, including the principle of proportionality, the principle of high-level protection as well as respect for fundamental rights enshrined in the Charter. Together, these principles require a demonstration of utilitarian welfare gain and ensure that costs to competing interests remain proportionate.

Investment incentives and engagement

As discussed in *Chapter 2.3*, the roadmap from idea to final image is lengthy and requires significant resources in terms of both skill and financial means. These costs are fixed and diminishing, meaning that the substantial parts are tied to the initial production. Once the system is ready to produce images, the marginal cost decreases substantially, and this cost is not significantly impacted by the number of images generated. These factors display characteristics consistent with the classical public good problem, creating potential for competitors to free ride and disincentivising production

However, it is important to acknowledge that the systems themselves often provide primary monetisation routes for developers through licensing, for example. Therefore, the additional copyright protection over individual images therefore represents a secondary layer of protection, the utilitarian justification of which must be evaluated separately. The possibility of exclusive rights over images is assumed to encourage creators to invest in and produce goods, which is particularly important where system-level monetisation is insufficient or impractical. This secondary protection may be most justified for specialised systems with limited commercial appeal, where the developer cannot recover costs through system licensing alone. Otherwise, the art of image generation risks market failure due to underproduction.

Although eligibility for copyright does not eliminate the aforementioned costs, it can significantly alter the risk/return equation, as exclusive rights provide an alternative monetisation strategy. If system-level profits are inadequate, licensing the images can motivate both a quantitative and qualitative increase in the images. This could be particularly relevant for developers of highly specialised systems, such as the previously mentioned photographer who wants to produce photorealistic images of imaginary animals. The images from this system would be so specialised that few other individuals would be interested in licensing the system itself, making image-level copyright protection necessary for financial viability.

Meanwhile, claims to authorship by developers could create barriers for the users who are involved in the generation. As the previous examination has hinted at, users can have a substantial impact on the images, which could qualify them for authorship. If they are forced into co-authorship, monetisation of the images they have initiated becomes more complicated. This could lead to reduced user engagement, thereby limiting the social benefits of the system overall.

Market structure and competition

While the recognition would increase the developers' chances of recouping their investment, the high entry barriers already favour a small group of well-capitalised entities. Should the same actors acquire enforceable rights over the images, it is possible that distortionary effects could lead to a deepened market concentration. The application of the proportionality principle requires careful assessment of whether identified risks of market concentration outweigh the claimed utilitarian benefits.

Firstly, it is important to note that the rights associated with the actual pictures do not hinder independent research. This ensures that there is ample space for subsequent innovation of the systems. Secondly, the possibility of legally protected returns has the potential to stimulate licensing markets that would otherwise be commercially irrational, thereby expanding access to smaller creative start-ups that opt to license the images rather than building models themselves. However, this potential benefit must be weighed against the enforcement barriers that may limit the ability of smaller developers to realise these rights in practice.

Volume and diversity of images

Predictably, a pure user-only allocation of authorship predictably channels developer focus into general-purpose models optimised for producing a large variety of general images, possibly leaving niche cultural domains underserved. However, the claim that developer authorship would incentivise specialised cultural content must be carefully qualified. While copyright protection could theoretically encourage the creation of micro-models fine-tuned to regional aesthetics, minority representation or idiosyncratic artistic visions, the high investment costs of culturally specific training data could mean that such specialisation is only accessible to well-capitalised entities. This could potentially limit rather than expand cultural diversity.

Nevertheless, where there are genuine market gaps for culturally specific content that cannot be served by general-purpose models, targeted copyright protection may provide the necessary incentives. The qualitative dividend here would be a richer, more pluralistic corpus of digital culture, provided that the

protection mechanism does not exacerbate existing barriers to cultural participation.

Access and freedom of expression

Exclusive rights create friction. For example, if the scope is not carefully limited, the volume of diverse images that can be produced risks flooding the market. Any model with a specialised style can quickly flood the market and effectively monopolise the field. Therefore, protection must be tailored to images that truly embody the developer's creative contribution rather than granting blanket entitlement to all future images. Furthermore, copyright always limits third parties' ability to use the works to express their own vision to a certain extent.

Developers already exercise significant de facto control over what is depicted in the images through filtering and model design. Adding another layer of control over images that they cannot create by themselves is in practice another way of potentially curbing the user's freedom of expression. However, if developers can maintain control as authors after the images are generated, this could reduce the risk of over-blocking during the generation.

Summary

The recognition of developers as the holders of copyright for AI-generated images within the EU legal framework, when assessed utilitarianly, reveals a complex balance of competing interests. While acknowledging developers as copyright holders could address the public good problem inherent in AI system development by providing secondary monetisation opportunities, particularly for specialised systems with limited commercial appeal, it simultaneously risks creating market concentration among well-capitalised entities and complicating user engagement through forced co-authorship arrangements. The analysis demonstrates that, while such recognition could theoretically incentivise culturally specific content and expand licensing markets for smaller entities, high barriers to specialised training data could paradoxically limit cultural diversity to well-funded actors. Moreover, the supplementary privilege pertaining to the ownership of generated images has the potential to impede the liberty of expression, despite the possibility of a reduction in over-blocking during the generation process. It is vital to acknowledge that any utilitarian justification for developer authorship must demonstrate proportionate welfare gains, and must do so through careful tailoring of protection to genuine creative contributions. Rather than granting blanket entitlement and ensuring compliance with EU principles of proportionality, high-level protection and fundamental rights, as they are enshrined in the Charter, the focus should be on more precise and targeted measures.

4.3 Labour theory

4.3.1 Introduction

Unlike utilitarianism, the labour theory justifies ownership and exclusive rights when the object in question is the result of an individual's labour.[153] This argument can be viewed from two perspectives: instrumental or normative. Either exclusive rights created through the creative process incentivise labour, or protection is justified because labour should be rewarded.[154] In any case, labour should be compensated.

In this context, the term *labour* has no universally accepted definition and should be interpreted in light of the actions accepted in the doctrine. Labour encompasses the act of an individual adding something to an unclaimed object, mixing it with something of their own, or improving it in another way.[155] This can be as simple as gathering fruit from an unclaimed tree, which justifies a gatherer's claim to the fruit over that of others.[156] However, it is a requisite that the labour does not fully exploit the resources to the point that no others cannot benefit from them. This can also be achieved by letting others enjoy a certain part or aspect of the produce. [157] This gives everyone the opportunity to gain as much as they are willing to work for. [158]

As with the concept of labour, which is challenging to define, the quantification of labour is similarly difficult. To elaborate further, the extent to which an action can be categorised as labour, and consequently the amount of remuneration that can be deemed appropriate, is a crucial consideration.[159] A key aspect of the reward is that it must be balanced in relation to the labour expended. This includes considerations of expenditure to competing interests caused by the labour.[160] In theory, one way to estimate the appropriate reward is to quantify the value that labour has brought to others' lives. [161] A more practical approach would be to use sales numbers or similar. However, these would need to be predicted in advance. In practice, a painter only knows the value of their painting once it has been sold, and perhaps not even then, but only after it has been sold a few times. Yet, for the first sale to happened, the painter must have a rightful claim to the piece to begin with.[162]

In the context of creative work, the labour refers to the effort that an artist puts into a piece by executing their ideas and inspiration, which untouched

[153] Rognstad (n 130) 30.
[154] ibid.
[155] Lukoševičienė (n 82) 193; Rognstad (n 130) 29.
[156] Ang (n 5) 81.
[157] ibid 83.
[158] Justin Hughes, 'The Philosophy of Intellectual Property' (1988) 77 Geo. L. J. 287, 8.
[159] Rognstad (n 130) 35.
[160] ibid 82; Lukoševičienė (n 82) 193.
[161] Hughes (n 154) 14.
[162] ibid 17.

are part of the common domain.[163] Monopolising these would endanger the principle of ensuring that there is enough for others to benefit from. [164] Elements of the labour theory are evident in both the statutory framework and the jurisprudence of EU copyright law. The DSM Directive preamble provides instrumental justifications, stating that the framework stimulates creativity, investment, and the production of new content.[165] The same argument can be found in the InfoSoc Directive.[166] Furthermore, the latter directive's preamble emphasises the need for authors to be appropriately rewarded for their work and highlights the requirement for legal protection of that work.[167] This could be interpreted as either a normative statement, viewing financial compensation as a prerequisite for creation, or a descriptive statement, indicating that authors' costs must be covered by the products of their work. [168]

Support can also be found in case law, as the CJEU has made several statements that aligns with the theory. Notably, in the Infopaq and Levola Hengelo cases, the Court stated that intellectual property rights stem from human effort. However, in the Football Dataco case, the Court noted that it is not sufficient to expend labour and skill. Consequently, the *sweat of the brow doctrine* was disregarded as part of EU copyright law. Furthermore, the author is granted exclusive rights even when the work has no value, contradicting the justification theory.

4.3.2 The developer scenario

This section will examine the extent to which developers of GENAI systems should be recognised as copyright holders for images generated by their models, based on the labour theory. As previously stated, the justification for protection is evident when the object is a product of the individual's labour. Furthermore, it is imperative to ensure that the reward, authorship claim, is balanced in relation to the labour that has been expended. The analysis will be divided into two sections. Firstly, the question arises as to whether the developer's input can be considered as labour. Secondly, the question of whether it is appropriate to compensate that labour with the reward of authorship must be posed.

Labour definition

The development process can be summarised as producing an algorithm that can process raw data to absorb patterns and consequently produce novel images based on these patterns. From an abstract perspective, the combination

[163] Lukoševičienė (n 82) 194; Ramalho (n 20) 22.
[164] Ang (n 5) 83.
[165] 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 [2019] OJ L130/92, recital 2.
[166] InfoSoc Directive recital 2 and 4.
[167] Ibid, recital 10.
[168] Rognstad (n 130) 81.

of the algorithm, the training and the developer's ideas results in a system that can generate the desired image. The developer contributes with effort in many different forms, from more conceptual to manual. Many of the steps in the process might constitute labour that deserves reward in its own right, given that reward should be granted based on whether the product adds value for another person. Even when the developer has merely finished the problem definition stage, the plans produced may be considered valuable in their own right to those who lack the skill to produce something similar.

It is in the post-development phase that the developer's claim is truly put to the test. Upon completion of their contributions, the outcome is a system rather than an image. The latter requires two things: a user who initiates the generation and a system that generates the image. As a consequence, the relationship between the developer and the final product is somewhat indirect. This does not necessarily negate the developer's possibility of authorship claim; however, it is a matter that warrants further discussion. Whilst the labour theory acknowledges the value of collaborative work and does not stipulate a requirement for a direct connection, it is important not to exaggerate the extent of the developer's impact. A parallel may be drawn with a digital brush creator who develops a revolutionary new painting tool with unique capabilities for creating specific visual effects. Whilst the brush inventor's fundamental labour facilitates wholly novel forms of artistic expression, it should be noted that they are not the artist behind every painting created with their brush. However, it is imperative to acknowledge the pivotal role that technical innovation plays in defining the distinctive character of each artwork. This innovation is predicated on specific engineering choices pertaining to the brush's design and functionality. It is imperative to establish a limit that is reasonable in accordance with the objective of the theory.

Despite their absence from direct involvement in the generative process, the impact of the developer remains a significant factor. It is important to distinguish between models such as *This person does not exist* and models that enable user agency to select style, object, and so on. In the initial scenario, the developer's contribution is not the primary catalyst for the generative process. However, their efforts constitute the majority of the total labour. However, this is not necessarily the case in the second instance. As was indicated in Chapter 2.4, the user may invest a significant amount of effort in the production process. This may be regarded as either a collaborative endeavour or the user's autonomy can render the developer's efforts indistinguishable. Nonetheless, the initial interpretation is a more suitable perspective in accordance with the labour theory, which would imply a form of joint labour. Conversely, the analysis should conclude with the assertion that the reward may be subject to adjustment. The developer's efforts must not be disregarded on the basis that they are not visible in the image.

Balancing the rewards

The second question posed by this analysis is whether the rights associated with authorship are proportionate and appropriate in relation to the labour invested by the developer. As noted in Chapter 4.3.1, a crucial aspect of reward under labour theory is that it must be proportional to the labour expended. While the development process undeniably requires substantial time, skill and effort, the outcome is the system itself, not the image. The proportionality analysis must consider whether system-level protection provides adequate compensation. From an instrumental perspective, copyright protection for the GENAI system may sufficiently incentivise investment and production of new systems. Developers can monetise through licensing or direct sales, which appears relatively fair from a financial standpoint. However, this approach becomes problematic for specialised systems with limited commercial appeal. Referring again to the photographer seeking to generate images of imaginary animals in a specific style, the resulting system holds minimal value for third parties, despite the substantial labour invested.

From a normative perspective, the question becomes more complex. Granting developers authorship over all generated images raises fundamental concerns about the principle of leaving enough for others. Such monopolisation could endanger the common pool of ideas and expressions by conferring control over vast ranges of potential creative works, precisely what labour theory seeks to avoid.

The critical issue lies in the causal connection between developer labour and final image. While images would not exist without the developer's labour, they equally cannot exist without user intervention. This severs the direct link between developer and image, maintaining only the connection to the system. Therefore, while significant labour has been expended, granting full authorship over generated images appears problematic under labour theory's requirements for proportionality and commons preservation. System-level protection may fulfil the instrumental criteria, though the normative claim to output value remains complex.

Summary

In summary, the goal is to answer the question: Under EU copyright law, to what extent should developers of GENAI systems be recognised as copyright holders for images generated by their models, based on labour theory? Consequently, the question that needs to be answered is if the developer’s effort constitutes labour that results in the image in question, and if that labour should be rewarded with authorship eligibility. It has been noted that the image would not exist without the labour of the developer, however it does not exist solely because of that effort. Without intervention from the user, there

is no image at all. This severs the link between the developer and the object in question, while leaving the connection to the system. Therefore, while significant labour has been expended, granting full authorship over the final image appears problematic under the labour theory's requirements for proportionality and the preservation of the commons. As long as the developer can get an appropriate reward for the GENAI system itself, the criteria may be fulfilled from an instrumental perspective, though the normative claim to the value of the images remains complex.

4.4 Personality theory

4.4.1 Introduction

The personality theory builds upon the idea that connections can be formed between individuals and property. As the connection is established, the object becomes an externalised part of the author's own free will, and ultimately their own being. The claim to that object is justified as the object essentially is regarded as a part of the person. [169] The damage that the object is subjected to also affects the person behind it by affiliation. Therefore, the object is a vulnerability to the individual.[170] Concurrently, it functions as a medium for the expressing one's personality. [171] It can thus be posited that the individual has personal motives that that goes beyond financial recoupment or gains. [172] It is important to note that not all property carries this type of connection, and there is a variance in strength amongst those that do. [173] The strength of the connection between an object and a person is directly proportional to that person's claim to the object. [174] However, the necessity for increased protection is debatable.[175]

The connection can be realised in a variety of ways. Through different interactions with the world around them, individuals appropriate the external world to their own mind.[176] It is not necessary to create an object from scratch or use it to form a personal connection with it. The establishment of a connection may be achieved through the act of marking it, taking physical control of it, or imposing a form upon it. This can be done provided that the circumstances are appropriate, for instance that the object is not claimed by someone

[169] ibid 35.
[170] Mira T Sundara Rajan (ed), *Moral Rights: Principles, Practice and New Technology* (Oxford University Press 2011) 9.
[171] Ramalho (n 20) 23.
[172] Sundara Rajan (n 166) 4.
[173] Rognstad (n 130) 38.
[174] Ang (n 5) 85.
[175] Hughes (n 154) 38.
[176] ibid 33.

else.[177] There is no need to incentivise production, as the core of every individual is characterised by their will and constant seeking of actuality and effectiveness. [178] Both are central to creation of the connection. [179]

Some mediums make the personality less apparent in the object. Generally speaking, however, all categories of objects can carry it.[180] At first glance, it may be more intuitive to detect personality in works such as poems, music and novellas, than in those of a more technical nature, such as computer software. However, one could argue that stating that the latter cannot carry personality is an oversimplification of the art form.[181] It is essential to consider the object from the artist's perspective. A laymen might miss particularities that connect the artist to the object. Often, subtle manifestations are only recognisable to people who are intimate with the craft, such as what moves characterise a certain chess player.[182] It can be especially challenging to find connections in large collaborations, where individual can have varying degrees of involvement and thus claims.[183] However, not all actions of will cumulate to a personal connection. There must be a certain degree of freedom for the individual to make decisions, even if they are unconscious. This eliminates all choices that are severely limited, and allow for only practical considerations, such as those of technical nature.[184] Combining effort, competence and innovation could result in a product that one could take pride in, which is a way of detecting the connection. The objective of production is less important than the product itself, which can be created for entertainment or technical reasons.[185]

The concept of personality in works can be found in both the legal texts and case law of EU copyright law. However, the term *personality* is not necessarily meant to be interpreted in the same way. According to Recital 16 of the Term Directive, a work (more specifically, a photograph) is considered original if it reflects the author's own personality.[186]. The link between originality and personality has also been discussed by the CJEU in key cases such as Infopaq, Painer and Levola Hengelo, which were mentioned in the previous chapter. In summary, case law shows that a work must carry the author's personal imprint to qualify for protection. Personality is imprinted onto a work when the author makes free and creative choices. Without these choices, the

[177] ibid 35.
[178] ibid 32.
[179] ibid 39.
[180] ibid.
[181] Rognstad (n 130) 39.
[182] Hughes (n 154) 42.
[183] Sundara Rajan (n 166) 23.
[184] Enrico Bonadio, Luke McDonagh and Plamen Dinev, 'Copyright in Eorks Created by Artificial Intelligence: Between Creativity and Investments' in Enrico Bonadio and Patrick Russell Goold (eds), *The Cambridge handbook of investment-driven intellectual property* (Cambridge University Press 2023) 77.
[185] Sundara Rajan (n 166) 299.
[186] Term Directive recital 16.

product is not considered original and cannot be protected; likewise, it is not the product of free will.

Although moral rights are not harmonised, they are prevalent in all jurisdictions bound by the Berne Convention. This type of right recognises that authors have interests beyond financial gain and that their personal connection to their work is worthy of protection.[187] Moral rights are not just about recognition, but also about the integrity of the work and, indirectly, the integrity of the author.[188] Therefore, by principle, these rights are inalienable.[189] This highlights the unbreakable nature of the connection between a personal object and an individual. In an extreme analysis, the proposition of selling all rights to a work of art could be considered a form of enslavement.[190]

4.4.2 The developer scenario

The question of whether developers of GENAI systems can be recognised as authors of images created by their models hinges on the extent to which the images can be seen as an externalisation of the developer's own free will and personality. The foundation of this analysis is the assertion that property rights emerge when individuals imbue an external object with their will and personality. In the context of GENAI, this prompts the consideration of whether and how the developer's personal choices and values become embedded in the images in question.

Actions

Developers are entrusted with significant creative discretion at various stages of system development. This includes making sophisticated design decisions regarding the use case of the model; selecting appropriate training data; and determining the user's possibilities for interaction. Whilst these choices arguably allow the developer the greatest scope of creative freedom, there are numerous others that are subject to only a limited number of technical constraints. Cumulative, these choices may create a reflection of the developer's aesthetic vision, technical philosophy, and a comprehensive understanding of the system's priorities. The distinctive personality of a developer, understood here as the unique perspective and prioritisation they bring to problem-solving, can be traced in these many small but significant choices.

Lack of causal connection

However, the developer's choices only indirectly affect the generated images. While these choices impact the final product, allowing users to make creative choices carries a high risk of disruption. While collaborations are recognised,

[187] Sundara Rajan (n 166) 4.
[188] ibid 5.
[189] Ramalho (n 20) 23.
[190] Lukoševičienė (n 82) 195.

they make it difficult to discern personal connections. This is particularly pertinent given that the user has the final say on when the creative vision is executed, and the black box, which makes it hard to backtrack on the system's decision-making rationales. It is reasonable to assume that the connection can be measured along a spectrum of strength rather than as a direct or indirect binary distinction. The extent to which a developer's personality is reflected in an image depends substantially on the model's configuration and the agency allowed to the user. For example, in systems that are highly fine-tuned for a specific image style and impose robust filtering according to the developer's vision, one could argue that traces of the developer's creative vision are present in every image. An illustrative example of this would be a photographer wanting to generate photorealistic images of imaginary animals in a style similar to their real-life photography and thus produces a system that will not present anything else to the user. The image will be similar to the developer's vision no matter the user's intervention, which could prove a personal connection between the developer and the generated images. In contrast, it is more difficult for the developer to form a strong personal connection to the images when the system can generate a wide range of images in various styles and depicting different subjects, as with DALL-E. While there may still be traces of the developer's choices, it is questionable whether these would amount to a level that justifies the exclusive rights that EU copyright entails.

It is easy to imagine scenarios in which the developer is unaware of the of a certain image, such as when the user has a licence for the GENAI system. It would seem counterintuitive for the developer to have a connection to an object that they are unaware of. However, the developer's lack of awareness of the image does not necessarily weaken the personality connection, as *Chapter 4.2.1* acknowledges: *Decisions might even be unconscious*. From a moral rights perspective, consideration must be given to the photographer in the previous example who is renowned worldwide for their distinctive style. Even if the developer is unaware of the specific images being produced, the association between them and the image may create legitimate interests in work integrity. This is because moral rights protect against alterations that could damage the author's honour or reputation, given that others are aware of the association.

Another fundamental question is whether generative processes, despite their statistical nature, can serve as a vehicle for personality expression. The image is based on a function of learned patterns from datasets that are recombined according to the model's development and user interactions. However, personality theory recognises that the connection between person and object 'can come to life in different ways' and does not require direct physical manipulation. The important question is whether the developer's will and creative vision have been sufficiently embedded in the system's structure to establish the type of personal connection that justifies protection. The statistical nature of

pattern recognition and recombination may parallel human creative processes more closely than initially apparent: humans also draw on learned patterns and past experiences when creating, synthesising existing knowledge into new forms through conscious and unconscious mental processes.

Summary

While developers exercise substantial creative discretion through design choices, data curation, and system configuration, potentially embedding their aesthetic vision and technical philosophy into the model's structure, the strength of the personal connection varies significantly based on system specificity and user agency. The causal relationship between developer choices and final images exists along a spectrum, with highly specialised, fine-tuned systems enabling stronger personality connections than broad, general-purpose models. The analysis is complicated by two factors. Firstly, the statistical nature of generative processes means that it is difficult to predict specific imagery. Secondly, developer unawareness of these images also complicates the analysis. However, personality theory's flexible approach to object-person relationships suggests that meaningful personal connections can emerge through embedded creative vision rather than direct manipulation. The determination of developers as authors within the framework of personality theory is contingent upon the extent to which their cumulative choices and creative philosophy have been sufficiently embedded within the system's architecture to justify the exclusive rights that copyright protection entails. This determination must be made on a case-by-case basis, taking into account the specific characteristics of each GENAI system.

5 Conclusions

5.1 Research question 1

Under EU copyright law, to what extent can developers of generative AI systems be recognised as authors for images generated by their models, considering existing originality requirements and the concept of 'author's own intellectual creation' established in CJEU jurisprudence?

Developers of generative AI systems can, in principle, satisfy the EU originality requirement by demonstrating *free and creative choices* at key stages of system design. Most notably in problem definition, training data curation, architecture selection, and the application of filtering or fine-tuning mechanisms. These choices are neither purely technical nor wholly dictated by functional constraints and thus may reflect the developer's personality in the sense recognised in Infopaq and Painer. However, the causal connection between those design choices and any individual image is impaired by the system's stochastic processes and user-driven inputs. Where user agency is high, the developer's personal imprint on a particular image weakens; by contrast, highly specialised or narrowly constrained systems channel outputs more directly toward the developer's intended aesthetic, strengthening the originality link.

Under EU law, expression requires that works be identifiable *with sufficient precision and objectivity*. GENAI images can meet this test, but the indirect, multi-phased generative chain necessitates a flexible approach akin to photography and software precedents. CJEU case law permits recognition of authorship where preparatory design choices persist through to final image, even if invisible to end users.

In summary, developers can be recognised as authors, provided they exercise meaningful creative discretion in designing system parameters that causally influence the generated images and maintain sufficient control over user interaction to make sure that their personality is reflected in the images. General-purpose models offering extensive user control pose greater challenges, whereas fine-tuned, domain-specific systems have a higher chance of aligning with the EU originality and expression standards.

5.2 Research question 2

Under EU copyright law, to what extent should developers of generative AI systems be recognised as authors for images generated by their models, based on classical justification theories (including utilitarian, labour and personality-based approaches) and the general principles of EU law?

The baseline is free market and copyright is a construct that can be justified on different basis. First, where it maximises net social welfare by preventing

under-production while remaining proportionate to users' freedoms, according to the utilitarian theory. This points to the fact that authorship is justified only when monetisation of the system cannot recoup costs and filtering keeps user agency limited in a way that limits their freedom of expression. Second, when the developer has expended labour and skill that are causally reflected in the image and the exclusive rights granted are proportionate with their contribution, leaving sufficient freedom for others, under the labour theory. Even though the labour is imperative to the existence of the imagery, sole authorship is disproportionate as the user provides indispensable labour as well that have a direct link with the images. Third, when the system's design channels the developer's creative choices so clearly into the output that the image embodies the developer's personal imprint and any distortion of the image would harm their personhood, under the personality theory. Consequently, authorship is only justified when the images reflect the developer's predefined style, excluding images from models that can produce a wider variety of styles. Ultimately, this differentiated approach aligns with the EU proportionality principle by granting protection only where it is necessary to spur genuine investment, thereby striking a fair balance between developers' innovation incentives and the public's freedom to create, access and share new cultural expressions.

Legislation and international instruments

Agreement on Trade-Related Aspects of Intellectual Property Rights (adopted 15 April 1994, entered into force 1 January 1995) 1869 UNTS 299

Berne Convention for the Protection of Literary and Artistic Works (adopted 9 September 1886, as revised at Paris 24 July 1971, and amended 28 September 1979)

Charter of Fundamental Rights of the European Union [2012] OJ C326/391

Consolidated version of the Treaty on the Functioning of the European Union [2012] OJ C326/47

Council Directive 93/98/EEC of 29 October 1993 harmonising the term of protection of copyright and certain related rights [1993] OJ L290/9

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 [2001] OJ L167/10

Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs [2009] OJ L111/16

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 [2019] OJ L130/92

Table of cases

Case name	Case number	ECLI
Bezpečnostní softwarová asociace – Svaz softwarové ochrany v Ministerstvo kultury	C-393/09	EU:C:2010:816
Cofemel – Sociedade de Vestuário SA v G-Star Raw CV	C-683/17	EU:C:2019:721
Eva-Maria Painer v Standard VerlagsGmbH and Others	C-145/10	EU:C:2011:798
Football Dataco Ltd and Others v Yahoo! UK Ltd and Others	C-403/08 & C-429/08	EU:C:2012:115
Funke Medien NRW GmbH v Bundesrepublik Deutschland	C-496/17	EU:C:2019:623
Infopaq International A/S v Danske Dagblades Forening	C-5/08	EU:C:2009:465
Johan Deckmyn and Vrijheidsfonds VZW v Helena Vandersteen and Others	C-201/13	EU:C:2014:2132
Kwantum Nederland BV and Kwantum België BV v Vitra Collections AG	C-227/23	EU:C:2024:844
Levola Hengelo BV v Smilde Foods BV	C-310/17	EU:C:2018:899
Martin Luksan v Petrus van der Let	C-277/10	EU:C:2012:65
Nintendo Co. Ltd v PC Box Srl and 9Net Sr	C-355/12	EU:C:2014:25
Kwantum Nederland BV and Kwantum België BV v Vitra Collections AG, Opinion of AG Campos Sánchez-Bordona	C-227/23 (Opinion)	
Pelham GmbH and Others v Hütter and Others	C-476/17	EU:C:2019:624
SAS Institute Inc. v World Programming Ltd	C-406/10	EU:C:2012:259
SI, Brompton Bicycle Ltd v Chedech/Get2Get	C-833/18	EU:C:2020:461
Spiegel Online GmbH v Volker Beck	C-516/17	EU:C:2019:625

Bibliography

Books

Abbott R and Rothman E, 'A Utilitarian Approach to Copyright Law and Generative Artificial Intelligence' in Guillaume Tusseau (ed), Research Handbook on Law and Utilitarianism (Edward Elgar Publishing 2024)

Ang S, The Moral Dimensions of Intellectual Property Rights (Edward Elgar 2013)

Bonadio E, McDonagh L and Dinev P, 'Copyright in Works Created by Artificial Intelligence: between Creativity and Investments' in Bonadio E and Goold PR (eds), The Cambridge Handbook of Investment-Driven Intellectual Property (Cambridge University Press 2023)

Kruse R and others, Computational Intelligence: A Methodological Introduction (Springer International Publishing 2022)

Largan C and Morris T, Qualitative Secondary Research: A Step-By-Step Guide (SAGE Publications Ltd 2019)

Levin M, 'The Cofemel revolution – originality, equality and neutrality' in Rosati E (ed), The Routledge Handbook of EU Copyright Law (Routledge 2021)

Lukoševičienė A, Authors and Creative Users: Addressing the Conceptual Challenges of Digital Creativity to EU Copyright Law from the Perspective of the Author (Lund University (Media-Tryck) 2021)

Martin-Prat M, 'Twenty Years After—the State of EU Copyright Law' in Hayleigh Bosher and Eleonora Rosati (eds), Developments and directions in intellectual property law: twenty years of the IPKat (Oxford University Press 2023)

McCorduck P, Aaron's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen (WH Freeman 1991)

Negnevitsky M, Artificial Intelligence: A Guide to Intelligent Systems (2. ed., [Nachdr.], Addison-Wesley 20)

Peczenik A, Juridikens Teori Och Metod: En Introduktion till Allmän Rättslära [Legal Theory and Method: An Introduction to General Jurisprudence] (Faculty of Law, Lund University 1995)

Pila J and Torremans P, European Intellectual Property Law (2nd edition, Oxford University Press 2019)

Romalho A, 'The competence and rationale of EU copyright harmonization' in Rosati E (ed), The Routledge Handbook of EU Copyright Law (Routledge 2021)

Ramalho A, Intellectual Property Protection for AI-Generated Creations: Europe, the United States, Australia and Japan (Routledge 2021)

Reichel J, 'EU-Rättslig Metod [EU Legal Method]' in Maria Nääv and Mauro Zamboni (eds), Juridisk metodlära [Legal methodology] (2nd edn, Studentlitteratur 2018)

Rendas T, 'Fundamental rights in EU copyright law' in Rosati E (ed), The Routledge Handbook of EU Copyright Law (Routledge 2021)

Rognstad O-A, Property Aspects of Intellectual Property (Cambridge University Press 2018)

Rosati E, Copyright and the Court of Justice of the European Union (First edition, Oxford University Press 2019)

Smits J and Borghuis T, 'Generative AI and Intellectual Property Rights' in Bart Custers and Eduard Fosch-Villaronga (eds), Law and Artificial Intelligence : Regulating AI and Applying AI in Legal Practice (TMC Asser Press 2022)

Sundara Rajan MT, Moral Rights: Principles, Practice and New Technology (Oxford University Press 2011)

van Gompel S, 'Creativity, Autonomy and Personal Touch: A Critical Appraisal of the CJEU's Originality Test for Copyright' in Mireille MM van Eechoud (ed), The Work of Authorship (Amsterdam University Press 2015)

Zamboni M, 'Rättsdogmatisk Metod [Legal Dogmatic Method]' in Maria Nääv and Mauro Zamboni (eds), Juridisk metodlära [Legal methodology] (2nd edn, Studentlitteratur 2018)

Journal articles

Awais M and others, 'Foundation Models Defining a New Era in Vision: A Survey and Outlook' (2025) 47 IEEE Transactions on Pattern Analysis and Machine Intelligence 2245

Bandi A, Adapa PVSR and Kuchi YEVPK, 'The Power of Generative AI: A Review of Requirements, Models, Input–Output Formats, Evaluation Metrics, and Challenges' (2023) 15 Future Internet 260

Catanzaro Z, 'Beyond Incentives: Copyright in the Age of Algorithmic Production' (2023) 13 NYU Journal of Intellectual Property & Entertainment Law 1

De Silva D and Alahakoon D, 'An Artificial Intelligence Life Cycle: From Conception to Production' (2022) 3 Patterns 100489

Hugenholtz PB and Quintais JP, 'Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output?' (2021) 52 IIC - International Review of Intellectual Property and Competition Law 1190

Hughes JL, 'The Philosophy of Intellectual Property' (1988-1989) 77 Geo LJ 287

Lucchi N, 'ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems' (2024) 15 European Journal of Risk Regulation 602

Militsyna K, 'Human Creative Contribution to AI-Based Output – One Just Can('t) Get Enough' (2023) 72 GRUR International 939

Priora G, 'Copyright Law and the Promotion of Scientific Networks: Some Reflections on the Rules on Co-Authorship in the EU' (2019) 9 Queen Mary Journal of Intellectual Property 217

Rosati E, 'Why Originality in Copyright Is Not and Should Not Be a Meaningless Requirement' (2018) 13 Journal of Intellectual Property Law & Practice 597

Semmelrock H and others, 'Reproducibility in Machine-learning-based Research: Overview, Barriers, and Drivers' (2025) 46 AI Magazine

Škiljić A, 'When Art Meets Technology or Vice Versa: Key Challenges at the Crossroads of AI-Generated Artworks and Copyright Law' (2021) 52 IIC - International Review of Intellectual Property and Competition Law 1338

Conference papers

Weisz JD and others, 'Design Principles for Generative AI Applications', Proceedings of the CHI Conference on Human Factors in Computing Systems (ACM 2024) https://dl.acm.org/doi/10.1145/3613904.3642466 accessed 9 July 2025

Official publications

European Commission. Directorate General for Communications Networks, Content and Technology., 'Trends and Developments in Artificial Intelligence :Challenges to the Intellectual Property Rights Framework : Final Report.' (Publications Office
2020) https://data.europa.eu/doi/10.2759/683128 accessed 7 January 2025

European Union Intellectual Property Office (EUIPO), 'The Development of Generative Artificial Intelligence from a Copyright Perspective' (European Union Intellectual Property Office 2025) https://euipo.europa.eu/tunnel-web/secure/webdav/guest/document_library/observatory/documents/reports/2025_GenAI_from_copyright_perspective/2025_GenAI_from_copyright_perspective_FullR_en.pdf

Websites

'Controlled Customization | Generative AI on Vertex AI' (Google Cloud) https://cloud.google.com/vertex-ai/generative-ai/docs/image/edit-controlled accessed 11 July 2025

'DALL·E 2' (OpenAI) https://openai.com/index/dall-e-2/ accessed 11 July 2025

'DALL·E 3' (OpenAI) https://openai.com/index/dall-e-3/ accessed 11 July 2025

'Design Multimodal Prompts | Generative AI on Vertex AI' (Google Cloud) https://cloud.google.com/vertex-ai/generative-ai/docs/multimodal/design-multimodal-prompts accessed 11 July 2025

'Midjourney' https://www.midjourney.com/website accessed 11 July 2025

'Prompt and Image Attribute Guide | Generative AI on Vertex AI' (Google Cloud) https://cloud.google.com/vertex-ai/generative-ai/docs/image/img-gen-prompt-guide accessed 11 July 2025

'Stable Diffusion Online' (Stable Diffusion) https://stablediffusion-web.com/ accessed 11 July 2025

'The Next Rembrandt: Recreating the Work of a Master with AI' (Microsoft News Centre Europe, 13 April 2016) https://news.microsoft.com/europe/features/next-rembrandt/ accessed 9 July 2025

'This Person Does Not Exist' https://thispersondoesnotexist.com/ accessed 9 July 2025