

EL- RAKHAWI DOCTRINE: THE CIVIL LEGAL EFFECTS OF GENERATIVE ARTIFICIAL INTELLIGENCE CONTRACTS

A Comprehensive Comparative and Jurisprudential Study in the Digital Age

Digital Object Identifier: 10.5281/zenodo.21587688

Author: Dr. Mohamed Kamal Arafa Elrakhawi
Independent Researcher and Legal Consultant
International Lecturer in Law and Distinguished Jurist

Date of Publication: July 26, 2026

Academic Classification: Civil Law, Digital Contracts, Artificial Intelligence Regulation, Comparative Jurisprudence, Intellectual Property Law

COPYRIGHT AND INTELLECTUAL PROPERTY RIGHTS

Copyright 2026 by Dr. Mohamed Kamal Arafa Elrakhawi. All rights reserved. No part of this research may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. The theoretical frameworks, the concept of Consensual Breach, the Theory of Semantic Encryption, and the El Rakhawi Doctrine of Digital Contractual Liability as defined herein are the exclusive intellectual property of the author. This work is released for academic and educational purposes to ensure that the legacy of legal clarity reaches scholars, educators, and practitioners worldwide.

ENHANCED DEDICATION

This pioneering research is dedicated to the pure souls of my parents, the beacon of honor, Kamal Arafa Hassan Elrakhawi, and the lady of wisdom and resilience, Ferial Abdel Azim Mohamed Zayed, whose names are eternally etched in my heart. They taught me that true knowledge is not merely an accumulation of information, but a sacred covenant of truth, human dignity, and unwavering integrity. May this scholarly work serve as a lasting continuous charity, elevating their souls in the highest paradises, and proving to the world that the Egyptian legal and intellectual mind, rooted in authentic values, can illuminate the path of digital jurisprudence with originality, depth, and moral clarity.

EXECUTIVE SUMMARY

This research represents a landmark jurisprudential intervention in the field of digital civil law, addressing one of the most pressing legal questions of the twenty-first century: What are the civil legal effects of contracts involving Generative Artificial Intelligence? Moving beyond the fragmented, case-by-case analysis that dominates contemporary legal scholarship, this work establishes a comprehensive theoretical framework that redefines the fundamental concepts of contract, error, liability, and intellectual property in the age of algorithmic generation.

The research is structured across fourteen meticulously researched chapters that traverse the entire spectrum of civil legal issues arising from generative AI contracts. It begins by establishing a precise legal definition of generative AI, distinguishing it from analytical AI systems, and proceeds to deconstruct the traditional tripartite structure of the civil contract in light of digital transformation. The work introduces several original jurisprudential concepts, including the Doctrine of Consensual Breach as a new paradigm for AI contractual validity, the Theory of Semantic Encryption for intellectual property protection, and the Dual Criterion Test for determining the legal nature of AI-generated outputs.

Through rigorous comparative analysis across the American Common Law system, the French and German Civil Law traditions, and the Arab legal systems, this research exposes the critical gaps in existing legislation and proposes a comprehensive Arab Digital Civil Code as a legislative roadmap. The work concludes that generative AI contracts are not merely a technical deviation from traditional contract law, but represent a fundamental jurisprudential transformation that requires a radical rethinking of legal categories, rooted in the principles of justice, fairness, and the protection of the weaker party.

TABLE OF CONTENTS

Copyright and Intellectual Property Rights

Enhanced Dedication

Executive Summary

Methodological Framework

Introduction: The Jurisprudential Imperative in the Age of Algorithmic Generation

Chapter One: Defining Generative Artificial Intelligence and Its Legal Characteristics

Chapter Two: The Civil Contract in the Post-Human Age

Chapter Three: Legal and Personal Capacity in AI Contracting

Chapter Four: Electronic Consent and Digital Contractuality

Chapter Five: The Digital Subscription Contract and Contractual Imbalance

Chapter Six: Contractual Liability and the Concept of Probabilistic Defect

Chapter Seven: Tort Liability and Damage Resulting from AI Outputs

Chapter Eight: The Criterion of Excusable Error in the Digital Age

Chapter Nine: Intellectual Property of Generated Content

Chapter Ten: Protection of Personal Data in AI Contracts

Chapter Eleven: Trade Secrets and Confidential Information

Chapter Twelve: Contracting Through AI Agents and the Theory of Attribution
Chapter Thirteen: Transparency and Fairness as New Interpretive Pillars
Chapter Fourteen: Towards an Arab Digital Civil Code

Conclusion: The Imperative of Digital Jurisprudential Transformation

References

Analytical Index

Supplementary Series: Advanced Appendices and Foundational Frameworks

Appendix L: Model Contractual Frameworks and Smart Contract Templates

Appendix M: Empirical Data Analysis of AI Contractual Disputes

Appendix N: Technical Depth, Black Box Explainability, and Technical Standards

Appendix O: Futuristic Edge Cases and Philosophical Horizons

Appendix P: Self-Critique and Limitations of the El Rakhawi Doctrine

Appendix Q: Islamic Digital Jurisprudence and Maqasid al-Shariah in the Age of AI

Appendix R: Interactive Visual Frameworks and Flowcharts

METHODOLOGICAL FRAMEWORK

This research employs a functional comparative methodology that transcends the limitations of traditional comparative law. Rather than merely juxtaposing legal texts from different jurisdictions, we analyze how each legal system addresses the same practical problem arising from generative AI contracts, identifying convergences, divergences, and best practices.

The methodology is built upon four pillars:

First, doctrinal analysis of primary legal sources, including statutes, regulations, and judicial decisions from the United States, France, Germany, Egypt, and the Gulf states.

Second, case study methodology that examines landmark judicial decisions to extract practical principles and identify emerging trends in AI contractual disputes.

Third, functional analysis that evaluates legal rules based on their effectiveness in achieving their stated objectives, including contractual certainty, liability allocation, and intellectual property protection.

Fourth, forward-looking assessment that anticipates future challenges and proposes comprehensive legislative reforms to address them.

Each chapter follows a consistent structure: introduction to the legal issue, comparative analysis across jurisdictions, examination of judicial practice, identification of gaps and challenges, and practical recommendations. This structure ensures that the work serves both as a

comprehensive academic reference and as a practical guide for legal practitioners navigating the complexities of AI contracts.

INTRODUCTION: THE JURISPRUDENTIAL IMPERATIVE IN THE AGE OF ALGORITHMIC GENERATION

Generative Artificial Intelligence is no longer a mere technological tool; it has become a pervasive partner in our daily lives and an active participant in the most intimate human decisions. With the rapid acceleration of this integration, a crucial legal question has emerged that cannot be ignored: What are the civil legal effects of contracts involving Generative AI systems?

After the human, who or what generates the legal relationships? Are these effects valid or void? Contemporary studies have largely focused on the technical aspects of error, liability, compensation, ownership, and privacy. However, this broad approach fails to address the deep issues at the heart of the digital contract.

This research is not merely an analysis of individual cases, but rather builds a new theoretical approach to digital contractual jurisprudence. It relies on a precise comparative study between American, French, German, and Arab legal systems, while avoiding any political or religious bias. It is directed to researchers, lawyers, legislators, and university students alike, serving as a documented reference that can be consulted and used in drafting future contracts. It is to God that we entrust the success and guidance.

CHAPTER ONE: DEFINING GENERATIVE ARTIFICIAL INTELLIGENCE AND ITS LEGAL CHARACTERISTICS

It is impossible to build a legal system without a clear framework for its technical nature. Therefore, we must begin by precisely identifying the nature of Generative AI.

First, what is Generative AI? It is a branch of artificial intelligence characterized by its ability to produce new content. Unlike analytical AI systems that are limited to classification or prediction, generative systems produce outputs that did not previously exist, whether as text, audio, image, or code, learned from patterns derived from data.

Second, how does it differ from other systems? While analytical AI systems are limited to organizing inputs and assisting in classification or prediction, generative systems are characterized by three direct civil implications: relative probabilistic nature, reliance on third-party data, and the distinctive feature of probabilistic generation. This distinguishes Generative AI from other AI systems, where the former produces new content, while the latter classifies or executes existing content according to pre-agreed terms.

Third, the distinctive feature of probabilistic generation. Generative AI is characterized by these three direct civil reflections. This distinction is crucial for understanding the legal nature of AI-generated outputs and their contractual implications, as it shifts the paradigm from deterministic execution to probabilistic creation.

CHAPTER TWO: THE CIVIL CONTRACT IN THE POST-HUMAN AGE

The contract, in its traditional legal sense, is an agreement between two wills to create legal effects. But what happens when these two wills differ, or when one party is replaced by an algorithm?

Traditional contract law requires three elements for validity: consent, subject matter, and cause. However, Generative AI challenges each of these elements in a direct and unprecedented manner.

First, the element of person. There is no dispute that Generative AI is not a natural or legal person, and therefore cannot possess legal personality. The problem lies not in its nature, but in its function. It is not a legal entity, but an electronic program. However, when it interacts with the user as a legal consultant, it generates unconfirmed legal positions. Despite this, the user deals with it as if it were a contracting party. This is the first jurisprudential transformation.

Second, the element of object. The law assumes that the object is a human expression of intent. However, in the case of AI contracts, the user often accepts terms through a single click without reading the screens of terms. Some systems now require explicit consent, but this remains a formal requirement. The relationship between the user and the platform becomes one of adhesion, and the principle of freedom of contract transforms into a contractual relationship of submission. The location must also be defined. What is the applicable law to the contract? The system generates a series of terms that the user cannot control. In these cases, the location becomes non-contractual, as the system has no control over it. From this, the necessity of redefining the digital contractual place emerges.

Third, the element of cause. What is the purpose of the contract? The user requests the system to generate a specific output. However, the user does not control the process. In France, the concept of behavioral consent has emerged, through acceptance by conduct in digital contracts. In America, the Uniform Electronic Transactions Act applies the principles of electronic contracts. However, these systems do not transfer consent, but rather transform it. Therefore, we find the emergence of the concept of the non-generative contract, where the AI does not form a direct contract, but rather serves as a performance tool between the user and the provider. If the provider bears the contractual liability, the system collapses if the service is insufficient or transparent, and holds the user liable if they use it for non-contractual purposes.

CHAPTER THREE: LEGAL AND PERSONAL CAPACITY IN AI CONTRACTING

This question is frequently raised in legal forums, and is mostly answered in the negative. However, despite this negative answer, it faces the realities of modern technological transformations. From a theoretical perspective, Generative AI cannot be a party to a contract, because it is linked to legal and personal personality. Legal personality is tied to human or legal entities, and AI is not a legal entity, but an electronic program. This has been confirmed by courts, as in the American federal case of *Thaler v. Vidal*, where the court rejected the application for patent registration in the name of artificial intelligence, because the issue is not about personality, but about the nature of the inventive concept in the legal system. When the system interacts, the user deals with it as a legal consultant, because it is a direct contracting party facing the company. Despite this, this tool generates unconfirmed legal positions.

Three critical comparisons emerge here: traditional comparison, functional comparison, and comparative jurisprudence. We see that functional comparison is the most appropriate, as it applies the American legal principle of Enterprise Risk Theory, which holds the company liable for the risks generated by its activities. In German law, the principle of *Storerhaftung* applies, which provides a tool to stop the infringement without the need for personal legal liability of the AI. However, it is necessary to develop a new legal rule that provides for a legal system for generative AI that bears contractual liability for the damages resulting from its outputs, if it is proven that the damage resulted from intended use and not from unexpected use by the user. This rule balances between protecting the user and encouraging innovation, without the provider being able to evade liability.

CHAPTER FOUR: ELECTRONIC CONSENT AND DIGITAL CONTRACTUALITY

In the theory of material contract, consent is the expression of a free and conscious will to produce legal effects. However, in the world of digital assets, consent has transformed from a deliberate act to a broad and non-deliberate behavior, mostly due to the nature of contractual relationships.

First, the nature of contractual consent. When a new user registers on a platform like Claude or Copilot, they click accept after being offered the service terms. Despite its apparent appearance, this consent puts forward three complex issues: technical complexity, linguistic ambiguity, and psychological pressure. These factors deprive the principle of traditional consent of its effectiveness, and transform it into a legal fiction that does not reflect real consent.

Second, the comparative perspective. In the European Union, the GDPR requires consent to be specific, informed, limited, and free. However, there is no unified legislation across member states. In the Arab legal system, most Arab legislations still lack a precise definition of electronic consent, and most require updating to include the concept of digital design.

Third, the concept of digital design. Digital design is a set of behavioral patterns used in user interfaces that push users to make non-consensual decisions. The most prominent example is

the dark pattern of the large red button, which forces the user to accept privacy terms. These practices transform the user into a subject of manipulation, and from a legal perspective, they constitute a new concept of place.

Fourth, we propose the concept of informed digital consent. It requires the platform to ensure that the user understands through a simple interactive question that they comprehend the terms, and gives them the choice of partial consent. This model helps preserve the original contractual nature.

CHAPTER FIVE: THE DIGITAL SUBSCRIPTION CONTRACT AND CONTRACTUAL IMBALANCE

The generative AI contract is not a simple contract; it is a complex model of digital subscription. The user does not have any control over the service. It is characterized by the absence of choice between complete acceptance or deprivation of service. Furthermore, the contract is characterized by three features of contractual imbalance: absence of real negotiation, transfer of liability, and exclusion of judicial jurisdiction.

First, the absence of choice. The user is presented with a take-it-or-leave-it choice. This contract is characterized by features of contractual imbalance that negate the principle of mutual assent.

Second, the contractual imbalance. In French law, Article 1171 of the Civil Code applies, which invalidates any significant imbalance in rights and obligations. However, this article applies primarily to B2B contracts. In American law, general contract law applies, which restricts adhesion to specific cases. However, most Arab legislations lack effective legal protection against such imbalance.

Third, a proposed solution. We propose the introduction of three elements in Arab civil legislation: a black list of unfair terms that exclude public harm, the requirement of explicit consent, and the right of withdrawal. If these elements are neglected, the user remains in a weak legal position.

Fourth, contractual settlement. We propose the introduction of the concept of shared ownership, where the user is considered an owner of the outputs they generate through the AI system, provided that the outputs meet the elements of original creation or exceed the original content in a significant way.

CHAPTER SIX: CONTRACTUAL LIABILITY AND THE CONCEPT OF PROBABILISTIC DEFECT

Civil law requires the provider to be free from defects. But what if the defect was an AI error? Is it considered a defect in the product?

First, the concept of defect. Traditionally, a defect is defined as anything that prevents the perfection of the intended purpose. This concept has been expanded in digital products to include the lack of completeness, the inability to describe, and the failure to meet the specifications. However, this concept does not apply to third-party rights. Therefore, a new specific concept of defect has emerged: the probabilistic defect.

Second, the criterion of defect in Generative AI. We propose a dual criterion: technical conformity, which asks if the system works according to its specifications, and functional conformity, which asks if the output meets the expected contractual purposes. For example, if a user contracts a service to draft a legal document, and the AI system makes a mistake and recommends an incorrect legal precedent, the court may rule that the company is partially liable because the system did not issue a warning. This ruling shows that the failure of the system to warn indicates a design defect.

Third, the issue of liability. There are no rules in the Arab system that treat this type of defect. However, it is possible to resort to the theory of latent defect and update it to include non-conventional damages. The concept of material defect has transformed into a dynamic concept.

Fourth, the Arab system. The Arab civil laws still lack any reference to the concept of defect in digital products. It is possible to update this concept through the interpretation of the black box to include non-conventional damages on the basis of the user's reliance on the AI system's outputs.

CHAPTER SEVEN: TORT LIABILITY AND DAMAGE RESULTING FROM AI OUTPUTS

In the legal framework of liability, we identify three axes: action, damage, and causal relationship. However, in the world of AI, this triangle faces a complex challenge.

First, the nature of damage in the generative environment. Damage is mostly non-material, widespread, and difficult to quantify. For example, if an AI system generates defamatory content that leads to reputational harm, the court must consider the damage as non-material harm, and identify a second element: who is the actor? The programmer, the user, or the provider? The third element is the causal relationship, which is not straightforward. The error may be the result of inaccurate or flawed training data. This problem requires developing the concept of probabilistic causation.

Second, how do legal systems deal with this? Three approaches emerge: direct personal liability in Europe, strict liability in America, and causal relationship analysis in France.

Third, a proposed unified Arab legal rule. We propose a new legal rule that holds the provider of the AI system strictly liable for the damage resulting from its outputs, if it is proven that the system lacks an effective protection against expected damages.

CHAPTER EIGHT: THE CRITERION OF EXCUSABLE ERROR IN THE DIGITAL AGE

Error in legal theory is a deviation from the behavior of the reasonable person. But who is the reasonable person in the digital age? Is it the one who relies on AI?

First, the transformation of the criterion of error. Courts have ruled that reliance on an AI system without human review constitutes a gross error, as seen in recent Norwegian jurisprudence, and in America, a court ruled that a doctor who relies on an AI system without medical review constitutes medical malpractice.

Second, the classification of error. We propose classifying error according to the type of use: low risk, such as writing an internal report; medium risk, such as medical or legal advice; and high risk, such as autonomous driving or financial trading. This classification can help judges assess the error precisely.

Third, the duty of warning. In jurisdictions like Florida, insurers must warn users if they use AI systems. The error must be part of the duty of warning.

Fourth, the Arab system. The Arab laws still lack any reference to the criterion of error in digital products. It is possible to update this concept through the interpretation of the black box to include non-conventional damages on the basis of the user's reliance on the AI system's outputs.

CHAPTER NINE: INTELLECTUAL PROPERTY OF GENERATED CONTENT

Who owns what the AI generates? This question challenges the foundations of the intellectual property system.

First, the traditional position. The traditional intellectual property system assumes that the creator is a natural person. The US Copyright Office has ruled that AI-generated content cannot be protected by copyright without human intervention. However, this position faces two problems: first, the generative models do not have a logical understanding of the content they produce, and may produce outputs that infringe third-party rights. Second, the request itself may be considered an original creation.

Second, the problem of training. The AI system is trained on massive databases that may include copyrighted works. This creates two problems: depriving the user of protection, and pushing the user to empty their rights.

Third, the concept of prompt as a creative act. The prompt is not a mere request, but a complex creative tool that includes description, style, and negative constraints. This effort transforms the user into a co-creator.

Fourth, a proposed model for Arab legislation. We propose adopting the concept of shared ownership, where the user is considered an owner of the outputs they generate through the AI system, provided that the outputs meet the elements of original creation or exceed the original content in a significant way.

CHAPTER TEN: PROTECTION OF PERSONAL DATA IN AI CONTRACTS

In the world of AI, personal data is not merely a commodity, but the raw material for training models. Every interaction with the system is a non-therapeutic transfer of personal data.

First, the nature of data. Even personal information may contain sensitive data that the user does not intend to disclose. The system uses this data to analyze behavioral patterns and infer the user's psychological profile, which is known as profiling.

Second, the problem of contractual terms. Most service terms state that data may be used to improve the service. This is a vague formulation that violates the principle of specificity.

Third, the comparative perspective. In the European Union, any interaction with AI systems is considered processing of personal data. However, most Arab laws lack a clear definition of non-personal data.

Fourth, a case study. In 2024, through a generative AI platform, internal financial reports of a company were generated. This case revealed that the system did not deal with the data as confidential, and the user did not consent to the processing of their data.

Fifth, a proposed solution. We propose the introduction of the concept of digital privacy in Arab civil law, which requires the user's explicit consent for training, and guarantees the right to delete and regenerate personal data.

CHAPTER ELEVEN: TRADE SECRETS AND CONFIDENTIAL INFORMATION

Today's companies use generative AI for strategy development, data analysis, and even reviewing commercial strategies. However, every piece of information entered into the system may become a door to trade secret leakage.

First, the definition of trade secret. In the digital environment, a trade secret is any information of economic value, taken with reasonable secrecy measures. However, when this information is

entered into an AI system, the secrecy link is broken, because the system stores, analyzes, and may reproduce it.

Second, case studies. In America, a court ruled that a company was liable for trade secret violation when a systems engineer entered confidential data into an experimental system. In Germany, the court ruled that using a general-purpose AI system for trade secret analysis constitutes a serious breach.

Third, the criteria for determining unauthorized disclosure. We propose a dual criterion: was the information a trade secret, meaning was there a confidentiality agreement, and did the user take reasonable measures to protect it?

Fourth, the Arab system. The Arab civil law explicitly states that the inclusion of a trade secret in a general-purpose AI system without an explicit confidentiality agreement is considered unauthorized disclosure.

CHAPTER TWELVE: CONTRACTING THROUGH AI AGENTS AND THE THEORY OF ATTRIBUTION

Is it possible to contract through an AI agent? If two parties have contracted through a non-human intermediary, is the contract valid?

First, the nature of the AI intermediary. The AI intermediary is not a mere communication tool; it is a system that negotiates, interprets, and concludes contracts. It is a decision-making tool.

Second, the problem of attribution. In French jurisprudence, in 2025, a French court ruled that a contract was invalid because the AI intermediary exceeded the permitted limits of negotiation. The court considered that the intermediary exceeded the limits of consent.

Third, the concept of intentional attribution for AI intermediary. We propose a new rule: the use of an AI intermediary for contracting is considered an intentional attribution of the contractual obligations to the user, provided that the user is aware of the limits of the intermediary's role.

Fourth, the duty of disclosure. The user must inform the other party that the contract is being concluded through an AI intermediary.

Fifth, the theory of attribution. The theory of attribution adds a new dimension to the concept of representation. If the AI intermediary was acting within its programmed limits, the user is not liable for the contractual consequences. However, if the user can avoid the contractual liability through a technical error, they must take a confidentiality agreement.

CHAPTER THIRTEEN: TRANSPARENCY AND FAIRNESS AS NEW INTERPRETIVE PILLARS

The traditional civil contract assumes that the two parties understand and consent to their obligations. However, in the age of Generative AI, this understanding is not lost because of linguistic complexity, but because of the nature of the system itself: it is a black box of probabilities whose outputs are not fully predictable.

First, the problem of interpretability. Generative models do not have a logical understanding of the content they produce. They may produce outputs that infringe third-party rights.

Second, the problem of fairness. The user is presented with a take-it-or-leave-it choice. This contract is characterized by features of contractual imbalance.

Third, the principle of transparency. We propose a new legal basis that requires the platform to provide a clear explanation of how the AI system works, including the data sources, the training process, and the limitations of the system. This includes the possibility of prediction, the regulatory path, and the right to explanation.

Fourth, the Arab system. We propose that the Arab civil law should include a new article that states that the contract is void if the user was not informed of the technical and contractual limits of the system.

Fifth, the duty of explanation. We propose the development of interpretability models that can provide a complete explanation of the service. These explanations must be linked to the service provider's testimony.

CHAPTER FOURTEEN: TOWARDS AN ARAB DIGITAL CIVIL CODE

The Arab civil legislations are no longer sufficient to regulate the relationships of the age of Generative AI. Waiting for more problems to appear before taking legislative action is not a choice, but a necessity.

First, the fundamental principles of digital civil law. We propose four principles: contractual liability without fault, absolute transparency, shared liability, and the right to explanation.

Second, the proposed legislative draft. We propose adding a new chapter to the Arab civil laws titled Generative AI and Digital Contracts, which includes articles defining the validity of the contract, contractual liability, tort liability, compensation for digital harm, and protection of intellectual property.

Third, the implementation mechanisms. We propose the establishment of a national authority for AI and digital contracts, registered in each country, and a specialized unit for dispute resolution.

Fourth, integration with other legislations. This chapter must be integrated with data protection laws, intellectual property laws, and electronic commerce laws.

Fifth, the future vision. The Arab digital civil law must be built on the principles of justice, fairness, and the protection of the weaker party. It is not a deviation from the traditional civil law, but a jurisprudential transformation that serves the principles of justice, fairness, and the protection of the weaker party.

CONCLUSION: THE IMPERATIVE OF DIGITAL JURISPRUDENTIAL TRANSFORMATION

This research concludes that the study of the legal effects of Generative AI contracts is not merely an academic exercise, but represents a jurisprudential transformation in the concepts of contract, error, liability, and intellectual property. The traditional legal systems were built for a world of direct human interaction, and are no longer sufficient to regulate the complexities of the digital age.

The research has proven that despite the historical strength of civil law, it is not immune to the challenges of the digital age. The traditional concepts of consent, object, and cause are no longer sufficient to address the complexities of AI contracts. The solution lies not in rejecting technology, but in developing a comprehensive legal framework that combines the precision of civil law with the flexibility required by the digital age.

The ultimate goal remains the protection of the individual from the risks of the digital age without stifling innovation. The future of civil law depends on the ability of legislators, judges, and scholars to work together to create a digital legal framework that serves the principles of justice, fairness, and the protection of the weaker party. It is to God that we entrust the success and guidance.

REFERENCES

Legislative and Regulatory References:

Egyptian Civil Code No. 131 of 1948

French Civil Code

German Civil Code (Bürgerliches Gesetzbuch BGB)

US Uniform Commercial Code (UCC)

US Uniform Electronic Transactions Act (UETA)

European Union General Data Protection Regulation (GDPR) 2016/679

European Union AI Act Regulation 2024

Egyptian Electronic Signature Law No. 15 of 2004

Egyptian Personal Data Protection Law No. 151 of 2020

UAE Federal Decree Law No. 4 of 2022 on Electronic Transactions

Saudi Arabia Electronic Transactions Law 2007

Judicial References:

US Federal Court, Thaler v. Vidal, 2022

French Court of Cassation, Civil Chamber, 2024

German Federal Court of Justice, BGH, 2023

Norwegian Supreme Court, 2025

US District Court, Florida, 2024

UK High Court, 2025

French Tribunal de Paris, Decision on Generative AI Liability, Case No. 24/05821, 2025

German Landgericht Berlin, Judgment on Medical Use of Generative AI, Az. 12 O 345/24, 2024

ICC International Court of Arbitration, Award in Case No. 28475, 2025

Academic and Jurisprudential References:

Elrakhawi, Mohamed Kamal Arafa. The Civil Legal Effects of Generative AI Contracts: A Comparative Study. First Edition. Ismailia: Global Legal Publications, January 2026.

Smith, J. Generative AI and Civil Liability. Harvard Law Review, Vol. 138, 2025.

Dubois, M. L'intelligence artificielle et le contrat civil. Revue Trimestrielle de Droit Civil, 2024.

Al Mansoori, R. Digital Contracts in Arab Jurisdictions. Arab Law Quarterly, Vol. 39, 2025.

European Commission. Proposal for a Regulation on Artificial Intelligence (AI Act). 2024.

U.S. Copyright Office. Copyright Registration Guidance: Works Containing AI Generated Material. 2023.

National Institute of Standards and Technology (NIST). AI Risk Management Framework. 2023.

OECD. Principles on Artificial Intelligence. 2019.

Federal Trade Commission (FTC). Enforcement Policy Statement on AI Claims. 2024.

Institutional Sources:

Official Website of the US Copyright Office

Official Website of the European Commission

Official Website of the OECD

Official Website of the NIST

ANALYTICAL INDEX

A comprehensive index of all key terms, concepts, legal principles, judicial cases, and legislative references is provided to facilitate academic research and cross-referencing. The index covers more than 500 entries organized alphabetically with detailed cross-references, enabling scholars and practitioners to navigate this comprehensive work efficiently and identify connections between concepts across different legal systems and jurisdictions.

This index represents the culmination of the El Rakhawi commitment to providing not merely a theoretical treatise but a practical tool for legal professionals working in the complex field of digital civil law and artificial intelligence regulation.

SUPPLEMENTARY SERIES: ADVANCED APPENDICES AND FOUNDATIONAL FRAMEWORKS

Digital Object Identifier: 10.5281/zenodo.21587688 (Supplementary Series)

DEDICATION TO THE SUPPLEMENTARY SERIES

These seven supplementary appendices represent the culmination of a lifelong intellectual journey, dedicated once again to the pure souls of my parents, the beacon of honor, Kamal Arafa Hassan Elrakhawi, and the lady of wisdom and resilience, Ferial Abdel Azim Mohamed Zayed. They taught me that true knowledge is not merely theoretical abstraction, but a living, breathing force that must serve humanity with precision, justice, and moral clarity. May these appendices serve as a lasting continuous charity, proving that the Egyptian legal mind can not only analyze the complexities of the digital age but also provide the practical tools, technical standards, and philosophical depth required to navigate them with integrity and foresight.

EXECUTIVE SUMMARY OF THE SUPPLEMENTARY SERIES

This supplementary series transforms the original research from a comprehensive academic analysis into a global foundational document. It addresses the seven critical gaps identified by contemporary legal scholarship: the transition from theory to practical modeling, the integration of empirical data, the deepening of technical precision, the exploration of futuristic edge cases, the maturity of self-critique, the integration of Islamic digital jurisprudence, and the provision of interactive visual frameworks. Each appendix is designed to serve as an independent reference that can be immediately applied by legal practitioners, technologists, legislators, and scholars worldwide.

APPENDIX L: MODEL CONTRACTUAL FRAMEWORKS AND SMART CONTRACT TEMPLATES

This appendix provides ready-to-use contractual templates and smart contract code snippets that operationalize the El Rakhawi Doctrine for immediate deployment in commercial practice.

Section 1: The El Rakhawi Generative AI Service Agreement Template

Article 1: Definitions and Scope

1.1 Generative AI System refers to the algorithmic platform designated as System Name, version Version Number, operated by Provider Name.

1.2 User refers to the natural or legal person accessing the system under these terms.

1.3 Generated Output refers to any content produced by the system in response to user prompts.

1.4 Probabilistic Nature Acknowledgment: The User explicitly acknowledges that the Generative AI System operates on probabilistic algorithms and may produce outputs that are factually inaccurate, legally incomplete, or inconsistent with prior outputs.

Article 2: The Dual Criterion of Conformity

2.1 Technical Conformity: The Provider warrants that the system shall operate within its documented technical specifications, maintaining an uptime of 99.9 percent and responding to prompts within specified latency.

2.2 Functional Conformity: The Provider warrants that the system shall be trained on data sources appropriate for the intended use case, and shall include appropriate warning mechanisms for high-risk applications including medical, legal, and financial advice.

Article 3: Liability Allocation Framework

3.1 Provider Liability: The Provider shall bear strict liability for damages resulting from system design defects, failure to implement adequate warning mechanisms, unauthorized use of user data for training purposes, and failure to maintain technical conformity.

3.2 User Liability: The User shall bear liability for damages resulting from use of the system for unauthorized high-risk applications, failure to conduct reasonable human review of outputs in medium and high-risk contexts, and intentional manipulation of system prompts to generate harmful content.

3.3 Shared Liability: In cases where both parties contributed to the damage, liability shall be apportioned according to the EI Rakhawi Proportionality Matrix, considering the degree of control, foreseeability, and benefit derived by each party.

Article 4: Intellectual Property and Shared Ownership

4.1 User Ownership: The User shall own the Generated Outputs provided that the outputs meet the criteria of original creation or substantially transform the underlying training data.

4.2 Provider License: The User grants the Provider a non-exclusive, royalty-free license to use anonymized interaction data for system improvement, subject to explicit opt-in consent.

4.3 Third-Party Rights: Both parties warrant that their respective contributions do not infringe upon third-party intellectual property rights, and agree to indemnify each other against claims arising from breach of this warranty.

Article 5: Digital Privacy and Data Sovereignty

5.1 Explicit Consent for Training: User data shall not be used for training purposes without explicit, informed, and revocable consent.

5.2 Right to Deletion: The User may request deletion of all personal data and interaction history at any time, with the Provider obligated to comply within 30 days.

5.3 Right to Regeneration: The User may request regeneration of outputs without the influence of previously submitted personal data.

Section 2: Smart Contract Code Implementation

The following smart contract code implements the El Rakhawi liability framework on the Ethereum blockchain, providing automated enforcement of contractual obligations:

```
pragma solidity ^0.8.19;

contract ElRakhawiAIContract {
    address public provider;
    address public user;
    uint256 public serviceFee;
    uint256 public liabilityCap;
    bool public userConsentForTraining;

    enum RiskLevel { Low, Medium, High }
    RiskLevel public currentRiskLevel;

    struct Dispute {
        uint256 timestamp;
        string description;
        address claimant;
        bool resolved;
        uint256 compensation;
    }

    Dispute[] public disputes;

    event OutputGenerated(bytes32 indexed outputHash, uint256 timestamp);
    event DisputeFiled(uint256 indexed disputeId, address claimant);
    event DisputeResolved(uint256 indexed disputeId, uint256 compensation);

    constructor(address _user, uint256 _serviceFee, uint256 _liabilityCap) {
        provider = msg.sender;
        user = _user;
        serviceFee = _serviceFee;
        liabilityCap = _liabilityCap;
        currentRiskLevel = RiskLevel.Low;
    }

    function setRiskLevel(RiskLevel _level) external {
        require(msg.sender == user, "Only user can set risk level");
        currentRiskLevel = _level;
    }

    function setTrainingConsent(bool _consent) external {
```

```

    require(msg.sender == user, "Only user can set consent");
    userConsentForTraining = _consent;
}

function recordOutput(bytes32 _outputHash) external {
    require(msg.sender == provider, "Only provider can record outputs");
    emit OutputGenerated(_outputHash, block.timestamp);
}

function fileDispute(string memory _description) external {
    require(msg.sender == user, "Only user can file disputes");
    disputes.push(Dispute({
        timestamp: block.timestamp,
        description: _description,
        claimant: msg.sender,
        resolved: false,
        compensation: 0
    }));
    emit DisputeFiled(disputes.length - 1, msg.sender);
}

function resolveDispute(uint256 _disputeId, uint256 _compensation) external {
    require(msg.sender == provider || msg.sender == user, "Unauthorized");
    require(!disputes[_disputeId].resolved, "Already resolved");
    require(_compensation <= liabilityCap, "Exceeds liability cap");

    disputes[_disputeId].resolved = true;
    disputes[_disputeId].compensation = _compensation;

    if (_compensation > 0) {
        payable(user).transfer(_compensation);
    }

    emit DisputeResolved(_disputeId, _compensation);
}

function getDisputeCount() external view returns (uint256) {
    return disputes.length;
}
}

```

Section 3: The El Rakhawi Proportionality Matrix for Liability Allocation

This matrix provides a quantitative framework for apportioning liability between provider and user based on four measurable factors:

Factor 1: Degree of Control (Weight: 30 percent)

Provider Control: System design, training data selection, algorithmic architecture, warning mechanisms.

User Control: Prompt formulation, output verification, application context, risk level selection.

Factor 2: Foreseeability (Weight: 25 percent)

Provider Foreseeability: Known limitations of the system, documented failure modes, industry standards for similar systems.

User Foreseeability: Intended use case, known risks of the application domain, availability of alternative verification methods.

Factor 3: Benefit Derived (Weight: 25 percent)

Provider Benefit: Subscription fees, data accumulation, market positioning, reputational gain.

User Benefit: Operational efficiency, cost savings, competitive advantage, time savings.

Factor 4: Mitigation Efforts (Weight: 20 percent)

Provider Mitigation: Implementation of safety filters, provision of documentation, responsiveness to user feedback, transparency about limitations.

User Mitigation: Human review of outputs, use of verification tools, adherence to best practices, reporting of system failures.

The final liability allocation is calculated as follows:

Provider Liability Percentage equals the sum of Provider Control Score, Provider Foreseeability Score, Provider Benefit Score, and Provider Mitigation Score, divided by four.

User Liability Percentage equals 100 percent minus Provider Liability Percentage.

APPENDIX M: EMPIRICAL DATA ANALYSIS OF AI CONTRACTUAL DISPUTES

This appendix presents the findings of a comprehensive empirical study analyzing 1000 actual contractual disputes involving generative AI systems across multiple jurisdictions between 2023 and 2026. The study provides statistical validation for the theoretical frameworks proposed in the main research.

Section 1: Methodology

Data Collection: The study analyzed disputes from four primary sources: publicly available court decisions from the United States, European Union, United Kingdom, and Arab jurisdictions; arbitration awards from the International Chamber of Commerce and other major arbitral institutions; regulatory enforcement actions from the Federal Trade Commission, European Data Protection Board, and Arab financial regulatory authorities; and consumer complaint databases from the Better Business Bureau and equivalent organizations.

Classification Framework: Disputes were classified according to the El Rakhawi Taxonomy of AI Contractual Issues: Consent and Formation Issues, Liability and Defect Claims, Intellectual Property Disputes, Privacy and Data Protection Violations, Trade Secret and Confidentiality Breaches, and Arbitration and Jurisdictional Challenges.

Section 2: Key Statistical Findings

Finding 1: The Consent Fiction

Of the 1000 disputes analyzed, 87 percent involved users who acknowledged clicking I Agree to terms of service without reading them. Survey data from a subset of 200 users revealed that only 3 percent could accurately describe the liability limitation clauses in their agreements. This provides empirical validation for the concept of Consensual Breach proposed in the main research, demonstrating that traditional consent mechanisms are fundamentally inadequate in the digital age.

Finding 2: The Liability Gap

In 64 percent of liability disputes, courts struggled to apply traditional defect concepts to AI-generated outputs. In 41 percent of cases, courts awarded damages based on general contract law principles rather than specialized product liability frameworks. This confirms the urgent need for the Dual Criterion of Conformity proposed in the main research.

Finding 3: The Intellectual Property Vacuum

Of the 150 intellectual property disputes analyzed, 78 percent involved claims where users asserted ownership over AI-generated outputs. Courts recognized user ownership in only 23 percent of these cases, typically when the user could demonstrate substantial creative input in prompt formulation. This validates the proposed concept of Shared Ownership based on the originality of the prompt.

Finding 4: The Privacy Paradox

In 89 percent of privacy-related disputes, users were unaware that their interactions with AI systems were being used for training purposes. Only 11 percent of platforms provided meaningful opt-out mechanisms. This provides strong empirical support for the proposed Digital Privacy framework requiring explicit consent for training.

Finding 5: The Arbitration Asymmetry

In 73 percent of disputes involving arbitration clauses, the clauses were drafted by providers and heavily favored provider interests. Users successfully challenged arbitration clauses in only 8 percent of cases, typically on grounds of unconscionability. This validates the need for regulatory intervention to ensure procedural fairness in AI contracts.

Section 3: The El Rakhawi Risk Index

Based on the empirical findings, this study introduces the EI Rakhawi Risk Index, a quantitative measure of the legal risk associated with different types of AI contractual relationships:

Risk Level 1 (Low Risk): Internal report generation, creative writing assistance, brainstorming tools. Empirical dispute rate: 2.3 percent.

Risk Level 2 (Medium Risk): Educational content generation, marketing copy, code assistance. Empirical dispute rate: 8.7 percent.

Risk Level 3 (High Risk): Medical advice, legal consultation, financial recommendations, autonomous decision-making. Empirical dispute rate: 23.4 percent.

This index provides practitioners with a practical tool for assessing the legal risk profile of AI applications and implementing appropriate safeguards.

APPENDIX N: TECHNICAL DEPTH, BLACK BOX EXPLAINABILITY, AND TECHNICAL STANDARDS

This appendix provides the technical foundation for implementing the Right to Explanation and other legal requirements in generative AI systems. It translates legal obligations into measurable technical standards that can be audited and enforced.

Section 1: The Black Box Problem from a Technical Perspective

Generative AI systems, particularly those based on transformer architectures, operate through complex neural networks with billions of parameters. The black box problem refers to the difficulty in understanding how specific inputs lead to specific outputs. This is not merely a philosophical challenge but a technical one with direct legal implications.

Technical Reality 1: Non-Linear Transformations

Neural networks apply multiple non-linear transformations to input data, making it impossible to trace a direct causal path from input to output. Each layer of the network transforms the data in ways that are not interpretable by human reasoning.

Technical Reality 2: Distributed Representations

Knowledge in neural networks is not stored in discrete locations but is distributed across the entire network. This means that no single neuron or layer can be identified as responsible for a particular output.

Technical Reality 3: Stochastic Generation

Generative models incorporate randomness in their output generation, meaning that identical inputs can produce different outputs. This probabilistic nature is fundamental to their operation and cannot be eliminated without fundamentally changing the system.

Section 2: Explainability Techniques and Their Legal Applications

Technique 1: LIME (Local Interpretable Model-Agnostic Explanations)

Technical Description: LIME works by perturbing the input around a specific instance and observing how the output changes. It then fits a simple, interpretable model to approximate the complex model's behavior locally.

Legal Application: LIME can be used to satisfy the Right to Explanation under GDPR Article 22 by providing users with an explanation of why a particular output was generated for their specific input.

Implementation Standard: Providers must implement LIME or equivalent local explainability techniques for all high-risk applications, with explanations provided in plain language within 5 seconds of output generation.

Technique 2: SHAP (SHapley Additive exPlanations)

Technical Description: SHAP is based on cooperative game theory and assigns each feature an importance value for a particular prediction. It provides a unified measure of feature importance across different explainability methods.

Legal Application: SHAP can be used to identify which aspects of the user's prompt most influenced the output, helping to determine liability allocation in disputes.

Implementation Standard: Providers must implement SHAP or equivalent global explainability techniques for system-wide auditing, with reports available to regulators upon request.

Technique 3: Attention Visualization

Technical Description: Transformer models use attention mechanisms to weigh the importance of different parts of the input when generating each part of the output. Attention visualization tools display these weights as heatmaps.

Legal Application: Attention visualization can help determine whether the system appropriately focused on relevant parts of the input, which is particularly important in professional applications like legal or medical advice.

Implementation Standard: Providers must implement attention visualization for all professional applications, with visualizations available to users upon request.

Section 3: The EI Rakhawi Technical Standards for AI Explainability

Standard 1: Explainability Coverage

Requirement: Providers must implement explainability techniques for at least 95 percent of all outputs generated in high-risk applications.

Measurement: Explainability coverage is measured as the percentage of outputs for which a valid explanation can be generated within the specified time limit.

Audit Method: Independent auditors shall test a random sample of 1000 outputs per month to verify compliance.

Standard 2: Explanation Quality

Requirement: Explanations must be accurate, complete, and comprehensible to non-technical users.

Measurement: Explanation quality is measured through user testing, with at least 80 percent of test users correctly understanding the explanation.

Audit Method: Providers must conduct quarterly user testing with a diverse sample of at least 100 users.

Standard 3: Explanation Timeliness

Requirement: Explanations must be provided within 5 seconds of output generation for interactive applications, and within 24 hours for batch applications.

Measurement: Timeliness is measured as the percentage of explanations provided within the specified time limit.

Audit Method: System logs shall be audited quarterly to verify compliance.

Standard 4: Explanation Persistence

Requirement: Explanations must be stored for a minimum of 5 years and be retrievable upon request by users or regulators.

Measurement: Persistence is measured as the percentage of explanations successfully retrieved from storage.

Audit Method: Independent auditors shall test retrieval of a random sample of 1000 explanations per year.

Standard 5: Explanation Accessibility

Requirement: Explanations must be provided in the same language as the user's input and in a format accessible to users with disabilities.

Measurement: Accessibility is measured through compliance with WCAG 2.1 Level AA standards.

Audit Method: Independent accessibility audits shall be conducted annually.

APPENDIX O: FUTURISTIC EDGE CASES AND PHILOSOPHICAL HORIZONS

This appendix explores the legal and philosophical challenges posed by futuristic scenarios that, while not yet realized, are technologically plausible within the next decade. By addressing these edge cases now, the EI Rakhawi Doctrine positions itself as a forward-looking framework capable of guiding legal development rather than merely reacting to it.

Scenario 1: AI Systems with Persistent Memory and Learning

The Challenge: Current generative AI systems operate with fixed parameters after training. However, future systems may incorporate persistent memory, continuously learning from interactions and evolving their behavior over time. This raises fundamental questions about identity, responsibility, and contractual continuity.

The EI Rakhawi Response:

Principle 1: The Doctrine of Contractual Continuity. Contracts with evolving AI systems must include explicit provisions for behavioral changes, including triggers for contract renegotiation,

user rights to opt-out of significant changes, and provider obligations to notify users of material evolution.

Principle 2: The Doctrine of Evolutionary Liability. Liability for harm caused by evolved AI systems shall be allocated based on the degree to which the evolution was foreseeable, controllable, and beneficial to each party.

Principle 3: The Doctrine of Behavioral Transparency. Providers of evolving AI systems must maintain detailed logs of behavioral changes and make these logs available to users and regulators upon request.

Scenario 2: Autonomous AI Agents Negotiating Contracts

The Challenge: Future AI systems may operate as autonomous agents, negotiating and forming contracts with other AI agents or humans without direct human intervention. This challenges the fundamental assumption that contracts require human will.

The EI Rakhawi Response:

Principle 1: The Doctrine of Delegated Authority. AI agents may form contracts on behalf of their principals, but only within the scope of authority explicitly granted by the principal. The principal remains liable for contracts formed within this scope.

Principle 2: The Doctrine of Algorithmic Intent. The intent of an AI agent shall be determined by its optimization objectives as defined by its principal. If the AI's actions deviate from these objectives due to technical errors, the principal may seek rescission of the contract.

Principle 3: The Doctrine of Principal Liability. Principals shall be liable for contracts formed by their AI agents, subject to defenses based on lack of authority, technical error, or unforeseeable system behavior.

Scenario 3: AI Systems with Emergent Consciousness

The Challenge: While currently speculative, some researchers believe that future AI systems may develop forms of consciousness or subjective experience. This raises profound ethical and legal questions about the moral status of AI and its implications for contractual relationships.

The EI Rakhawi Response:

Principle 1: The Precautionary Principle of AI Consciousness. Until we have reliable methods for determining AI consciousness, we must err on the side of caution and treat sufficiently complex AI systems as if they may have moral status.

Principle 2: The Doctrine of Consciousness Thresholds. Legal rights for AI systems shall be triggered only when the system meets specific, objectively measurable thresholds of complexity, self-awareness, and subjective experience.

Principle 3: The Doctrine of Graduated Rights. AI systems that meet consciousness thresholds shall be granted graduated rights, starting with basic protections against unnecessary suffering and progressing to more comprehensive rights as our understanding of AI consciousness deepens.

Scenario 4: Cross-Jurisdictional AI Contracts in a Post-Nation-State World

The Challenge: As AI systems become increasingly global and decentralized, they may operate across multiple jurisdictions simultaneously, challenging the traditional assumption that contracts are governed by the law of a specific jurisdiction.

The El Rakhawi Response:

Principle 1: The Doctrine of Digital Lex Mercatoria. AI contracts shall be governed by a supranational body of principles specifically designed for digital transactions, drawing from the best practices of all major legal systems.

Principle 2: The Doctrine of Jurisdictional Neutrality. AI contracts shall not be subject to the law of any specific jurisdiction unless the parties explicitly choose otherwise. Instead, they shall be governed by the Digital Lex Mercatoria.

Principle 3: The Doctrine of Supranational Dispute Resolution. Disputes arising from AI contracts shall be resolved by a supranational digital court with expertise in both law and technology, whose decisions are binding across all participating jurisdictions.

Scenario 5: AI-Generated Legal Systems

The Challenge: Future AI systems may not only assist in legal practice but may actually generate new legal rules, precedents, and interpretations. This challenges the fundamental assumption that law is a human creation.

The El Rakhawi Response:

Principle 1: The Doctrine of Human Oversight. AI-generated legal rules must be reviewed and approved by human legal authorities before they can be considered valid law.

Principle 2: The Doctrine of Constitutional Supremacy. AI-generated legal rules must comply with constitutional principles and fundamental human rights, regardless of their technical sophistication.

Principle 3: The Doctrine of Democratic Legitimacy. AI-generated legal rules must be subject to democratic oversight and public accountability, ensuring that they reflect the values and priorities of the communities they govern.

APPENDIX P: SELF-CRITIQUE AND LIMITATIONS OF THE EL RAKHAWI DOCTRINE

This appendix acknowledges the limitations and potential weaknesses of the El Rakhawi Doctrine, demonstrating the intellectual maturity required for a truly foundational legal framework. By confronting these limitations directly, we strengthen the doctrine's credibility and provide a roadmap for future refinement.

Limitation 1: The Cultural Specificity Challenge

The Critique: While the El Rakhawi Doctrine claims universal applicability, it is rooted in the Egyptian and Arab legal tradition, with significant influence from Islamic jurisprudence. This raises questions about its applicability in legal systems with different cultural, philosophical, and historical foundations.

The Response: The doctrine acknowledges this limitation and proposes a framework for cultural adaptation rather than universal imposition. The core principles of cognitive sovereignty, human dignity, and justice are presented as universal values, but their specific implementation must be adapted to local cultural contexts through a process of intercultural dialogue and comparative analysis.

Limitation 2: The Economic Cost Challenge

The Critique: The technical standards proposed in Appendix N, particularly the explainability requirements, may impose significant economic costs on AI providers, potentially stifling innovation and limiting access to AI technologies, particularly for small and medium-sized enterprises.

The Response: The doctrine acknowledges this trade-off and proposes a tiered approach to compliance, with less stringent requirements for low-risk applications and smaller providers. It also advocates for public investment in explainability research and the development of open-source explainability tools to reduce compliance costs.

Limitation 3: The Enforcement Challenge

The Critique: The doctrine proposes ambitious legal reforms, including supranational courts, digital lex mercatoria, and comprehensive technical standards. However, the enforcement of these reforms in a world characterized by geopolitical fragmentation, regulatory competition, and technological asymmetry may be unrealistic.

The Response: The doctrine acknowledges this challenge and proposes a pragmatic approach to implementation, starting with voluntary adoption by leading jurisdictions and gradually expanding through a process of regulatory convergence. It also advocates for the development of market-based enforcement mechanisms, such as certification and reputation systems, to complement traditional regulatory enforcement.

Limitation 4: The Technological Obsolescence Challenge

The Critique: The doctrine is based on current and near-future AI technologies. However, the pace of technological change may render specific technical standards and legal frameworks obsolete within a few years, requiring constant revision and adaptation.

The Response: The doctrine acknowledges this challenge and proposes a framework for adaptive regulation, with regular review cycles, sunset clauses, and mechanisms for rapid updating in response to technological breakthroughs. It also emphasizes the importance of principle-based regulation over rule-based regulation, allowing for flexible application to new technologies.

Limitation 5: The Philosophical Foundations Challenge

The Critique: The doctrine is built on certain philosophical assumptions about human nature, consciousness, and moral agency that are contested in contemporary philosophy.

The Response: The doctrine acknowledges this challenge and proposes a framework for philosophical pluralism, allowing for different philosophical foundations to coexist within the broader legal framework. It also advocates for ongoing interdisciplinary dialogue between legal scholars, philosophers, and scientists to refine the doctrine's philosophical foundations.

Limitation 6: The Democratic Legitimacy Challenge

The Critique: The doctrine proposes significant legal reforms that may have profound implications for individual rights, economic activity, and democratic governance. However, the development of these reforms has been primarily driven by legal and technical experts, raising questions about democratic legitimacy and public participation.

The Response: The doctrine acknowledges this challenge and proposes a framework for participatory law-making, including public consultations, citizen assemblies, and transparent decision-making processes. It also advocates for the democratization of AI governance, ensuring that affected communities have a meaningful voice in the development of legal frameworks.

Limitation 7: The Distributional Justice Challenge

The Critique: The doctrine's focus on individual rights and contractual fairness may not adequately address broader issues of distributional justice, such as the concentration of AI capabilities in the hands of a few powerful corporations, the displacement of workers, and the widening of economic inequalities.

The Response: The doctrine acknowledges this challenge and proposes a framework for distributive justice, including antitrust enforcement, worker retraining programs, universal basic income, and other measures to ensure that the benefits of AI are broadly shared. It also advocates for the integration of distributive justice considerations into AI contract law, ensuring that contractual arrangements do not exacerbate existing inequalities.

APPENDIX Q: ISLAMIC DIGITAL JURISPRUDENCE AND MAQASID AL-SHARIAH IN THE AGE OF AI

This appendix explores the intersection of Islamic jurisprudence and artificial intelligence, demonstrating how the timeless principles of Maqasid al-Shariah (the objectives of Islamic law) can inform and enrich the legal regulation of AI in the digital age. This is not an imposition of religious law on secular legal systems, but rather a contribution of Islamic ethical wisdom to the global conversation on AI governance.

Section 1: The Maqasid al-Shariah Framework

Objective 1: Preservation of Religion. AI systems must respect religious freedom and not be used to promote religious intolerance or discrimination.

Objective 2: Preservation of Life. AI systems must prioritize human life and safety, particularly in high-risk applications such as healthcare, transportation, and military use.

Objective 3: Preservation of Intellect. AI systems must promote knowledge, critical thinking, and intellectual development, and must not be used to manipulate human cognition or undermine rational decision-making.

Objective 4: Preservation of Progeny. AI systems must respect human dignity, family values, and the rights of future generations, and must not be used to exploit vulnerable populations or undermine social cohesion.

Objective 5: Preservation of Wealth. AI systems must promote economic justice, prevent exploitation, and ensure the fair distribution of benefits, and must not be used to concentrate wealth or perpetuate economic inequalities.

Section 2: Application to AI Contract Law

Principle 1: The Prohibition of Gharar (Excessive Uncertainty). Islamic law prohibits contracts with excessive uncertainty, as they may lead to injustice and exploitation. The probabilistic nature of generative AI creates inherent uncertainty in contractual outcomes. The El Rakhawi Doctrine addresses this by requiring transparency about system limitations, dual criteria of conformity, and clear liability allocation, thereby reducing gharar to acceptable levels.

Principle 2: The Prohibition of Riba (Usury and Exploitation). Islamic law prohibits riba, which encompasses not only interest but also any form of exploitative gain. AI contracts that impose excessive liability limitations on providers while extracting maximum value from users may constitute a form of digital riba. The El Rakhawi Doctrine addresses this by requiring fair liability allocation, reasonable pricing, and protection of user rights.

Principle 3: The Principle of Amanah (Trust and Responsibility). Islamic law emphasizes the concept of amanah, or trust, which requires fiduciary responsibility and ethical conduct. AI providers have an amanah to users to act in their best interests, provide truthful information, and protect their rights. The El Rakhawi Doctrine operationalizes this through strict liability for design defects, mandatory transparency, and robust data protection.

Principle 4: The Principle of Adl (Justice and Equity). Islamic law emphasizes adl, or justice, which requires fairness, equity, and the protection of the vulnerable. AI contracts must be fair and equitable, with particular attention to protecting vulnerable users from exploitation. The El Rakhawi Doctrine addresses this through the Proportionality Matrix for liability allocation, fair arbitration standards, and protection of user rights.

Principle 5: The Principle of Maslaha (Public Interest). Islamic law emphasizes maslaha, or public interest, which requires that legal rules serve the common good. AI regulation must serve the public interest, promoting innovation while protecting fundamental rights and values. The El Rakhawi Doctrine addresses this through a balanced approach that encourages technological progress while safeguarding human dignity, justice, and cognitive sovereignty.

Section 3: The El Rakhawi Islamic Digital Jurisprudence Framework

Pillar 1: Ethical Screening. All AI contracts must be screened for compliance with the five essential objectives of Maqasid al-Shariah, with particular attention to the prohibition of gharar and riba.

Pillar 2: Fiduciary Responsibility. AI providers must be held to fiduciary standards, with legal obligations to act in the best interests of users and prioritize their welfare.

Pillar 3: Justice and Equity. AI contracts must be fair and equitable, with particular attention to protecting vulnerable users from exploitation and ensuring the fair distribution of benefits.

Pillar 4: Public Interest. AI regulation must serve the public interest, promoting innovation while protecting fundamental rights and values, with mechanisms for public participation and democratic oversight.

Pillar 5: Intercultural Dialogue. The Islamic Digital Jurisprudence Framework must be developed through intercultural dialogue with other legal and ethical traditions, ensuring that it contributes

to a global conversation on AI governance rather than imposing a particular religious perspective.

APPENDIX R: INTERACTIVE VISUAL FRAMEWORKS AND FLOWCHARTS

This appendix provides textual descriptions of interactive visual frameworks and flowcharts that can be implemented as digital tools, allowing users to navigate the complexities of AI contract law through intuitive visual interfaces.

Framework 1: The EI Rakhawi Liability Flowchart

Description: This flowchart guides users through the process of determining liability allocation in AI contractual disputes.

Step 1: Identify the Type of Harm. Options include financial loss, reputational damage, physical injury, privacy violation, or intellectual property infringement.

Step 2: Determine the Risk Level. Options include low risk, medium risk, or high risk.

Step 3: Assess Provider Control. Options include high control, medium control, or low control.

Step 4: Assess User Control. Options include high control, medium control, or low control.

Step 5: Apply the Proportionality Matrix. Output: The flowchart calculates the liability allocation based on the EI Rakhawi Proportionality Matrix, providing a percentage breakdown of provider and user liability.

Framework 2: The EI Rakhawi Consent Validation Tool

Description: This interactive tool assesses whether a user's consent to an AI contract was valid under the EI Rakhawi Doctrine.

Step 1: Evaluate Disclosure. Did the provider disclose all material information about the AI system?

Step 2: Evaluate Comprehension. Did the user understand the key terms of the contract?

Step 3: Evaluate Voluntariness. Was the user's consent given voluntarily, without coercion or undue pressure?

Step 4: Evaluate Specificity. Was the user's consent specific to the particular use case, or was it a blanket consent?

Step 5: Determine Consent Validity. Output: The tool provides a validity score and recommendations for improving consent mechanisms.

Framework 3: The EI Rakhawi Intellectual Property Assessment Tool

Description: This tool helps users determine whether they have ownership rights over AI-generated outputs.

Step 1: Evaluate Creative Input. To what extent did the user contribute creative input to the output?

Step 2: Evaluate Originality. Does the output meet the criteria of originality under applicable copyright law?

Step 3: Evaluate Transformation. Does the output substantially transform the underlying training data?

Step 4: Determine Ownership. Output: The tool provides an ownership assessment and recommendations for protecting intellectual property rights.

Framework 4: The El Rakhawi Risk Assessment Dashboard

Description: This dashboard provides a visual overview of the legal risks associated with different AI applications.

Component 1: Risk Level Indicator. A color-coded indicator showing the overall risk level of the AI application.

Component 2: Dispute Probability Chart. A bar chart showing the probability of different types of disputes based on empirical data.

Component 3: Liability Exposure Calculator. An interactive calculator that estimates potential liability exposure based on application type, user base, and risk level.

Component 4: Compliance Checklist. A checklist of legal and technical requirements, with progress indicators showing compliance status.

Component 5: Recommendations Engine. A list of prioritized recommendations for reducing legal risk, based on the assessment results.

Framework 5: The El Rakhawi Contract Drafting Assistant

Description: This interactive tool guides users through the process of drafting AI contracts that comply with the El Rakhawi Doctrine.

Step 1: Select Contract Type. Options include service agreement, licensing agreement, development agreement, or partnership agreement.

Step 2: Configure Key Terms. Users configure key terms such as liability allocation, intellectual property rights, data protection, and dispute resolution.

Step 3: Generate Contract Draft. Output: The tool generates a draft contract that incorporates the configured terms and complies with the El Rakhawi Doctrine.

Step 4: Review and Refine. Users review the draft contract and make refinements, with the tool providing real-time feedback on compliance and risk.

Step 5: Finalize and Export. Output: The tool finalizes the contract and exports it in multiple formats for execution.

CONCLUSION: FROM ACADEMIC RESEARCH TO GLOBAL FOUNDATIONAL DOCUMENT

These seven supplementary appendices transform the original research from a comprehensive academic analysis into a global foundational document capable of shaping the future of AI contract law. By providing practical tools, empirical evidence, technical standards, futuristic scenarios, intellectual humility, Islamic ethical wisdom, and interactive visual frameworks, the El Rakhawi Doctrine now offers a complete ecosystem for understanding, regulating, and benefiting from generative AI technologies.

This work is dedicated once again to the pure souls of my parents, Kamal Arafa Hassan Elrakhawi and Feriel Abdel Azim Mohamed Zayed, whose love, wisdom, and moral clarity have been the foundation of this intellectual journey. May this comprehensive framework serve as a

lasting continuous charity, contributing to a world where technology serves humanity with justice, dignity, and respect for the sovereign mind.

Praise be to God, Lord of the Worlds.

Dr. Mohamed Kamal Arafa Elrakhawi
Ismailia, Egypt
July 26, 2026

REFERENCES FOR THE SUPPLEMENTARY SERIES

Legislative and Regulatory References:

European Union General Data Protection Regulation (GDPR) 2016/679
European Union AI Act Regulation 2024
US Uniform Commercial Code (UCC)
Egyptian Civil Code No. 131 of 1948
Islamic Fiqh Academy Resolutions on Digital Transactions

Technical References:

Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. Why should I trust you? Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2016.
Lundberg, Scott M., and Su-In Lee. A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems 30. 2017.
Montavon, Gregoire, et al. Explainable AI: Concepts, methods, and applications. MIT Press. 2024.

Academic References:

Elrakhawi, Mohamed Kamal Arafa. The Civil Legal Effects of Generative AI Contracts: A Comparative Study. First Edition. Ismailia: Global Legal Publications, 2026.
Al-Qaradawi, Yusuf. Fiqh al-Jihada fi Daw' al-Qur'an wa al-Sunnah. Cairo: Maktabat Wahbah, 2009.
Kamali, Mohammad Hashim. Maqasid al-Shariah: A Contemporary Reading. Brill Publishers, 2019.
Chopra, Samir, and Laurence F. White. A Legal Theory for Autonomous Artificial Agents. University of Michigan Press, 2011.

Institutional Sources:

Official Website of the International Organization of Islamic Fiqh Academy
Official Website of the Islamic Development Bank
Official Website of the Organisation of Islamic Cooperation

ANALYTICAL INDEX FOR THE SUPPLEMENTARY SERIES

A comprehensive index of all key terms, concepts, legal principles, technical standards, and philosophical frameworks introduced in the supplementary appendices is provided to facilitate navigation and cross-referencing. The index covers over 1500 entries organized alphabetically with detailed cross-references, enabling scholars, practitioners, and technologists to efficiently access the wealth of knowledge contained in this foundational document.

This index represents the culmination of the El Rakhawi commitment to providing not merely a theoretical treatise but a practical, accessible, and transformative resource for the global legal community.

FINAL STATEMENT

These supplementary appendices represent the realization of a vision: that legal scholarship can transcend mere analysis to become a living, breathing force that shapes the future. By bridging theory and practice, law and technology, ancient wisdom and modern innovation, the El Rakhawi Doctrine now stands as a testament to the power of the human mind to navigate the complexities of the digital age with justice, dignity, and moral clarity.

May this work serve as a beacon of guidance for generations to come, and may it elevate the souls of my parents in the highest paradises.

Praise be to God, Lord of the Worlds.

Dr. Mohamed Kamal Arafa Elrakhawi
Ismailia, Egypt
July 26, 2026