2025

DIGITAL TRANSFORMATION OF THE LEGAL PROFESSION IN THE AI ERA:

THE GEORGIAN CONTEXT

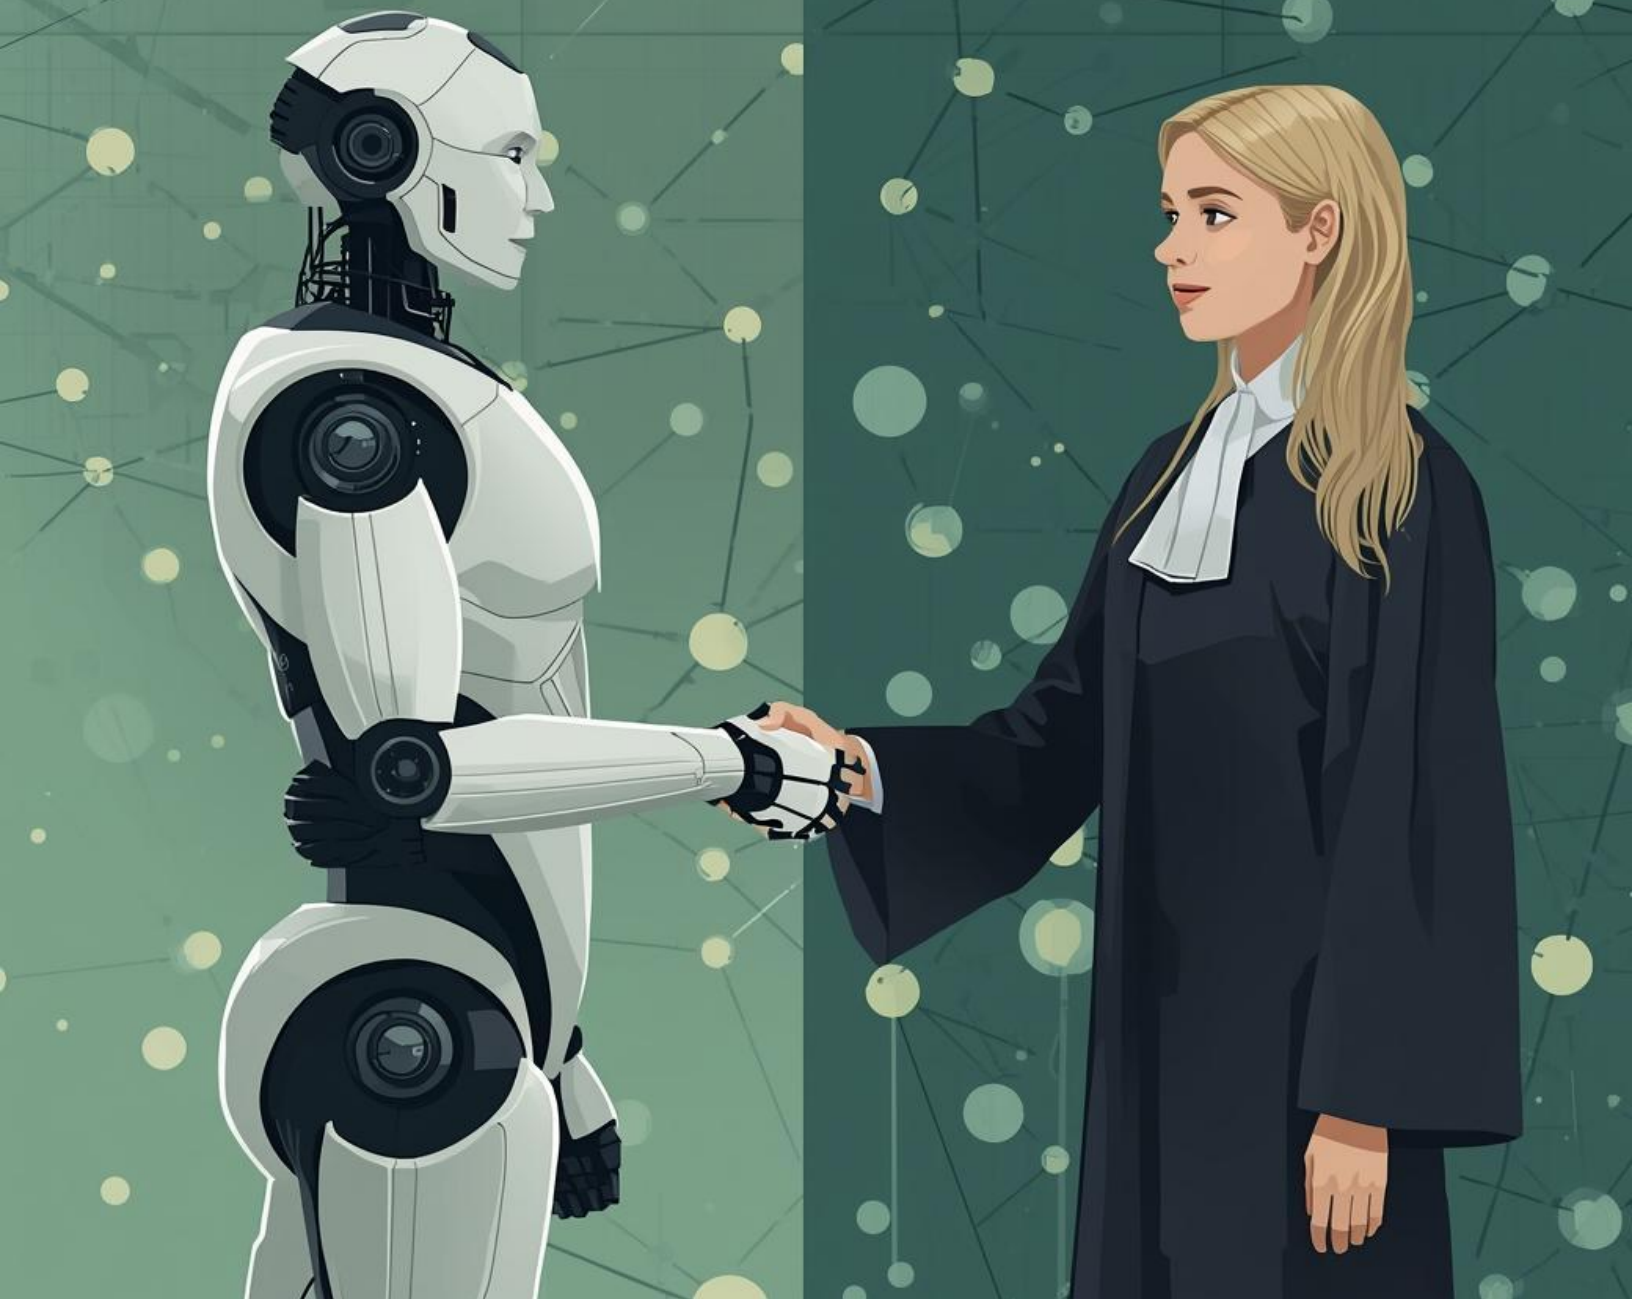

Acknowledgements

The project team express their sincere gratitude to all those who contributed to the development of this report and made this initiative possible.

First and foremost, we extend our deepest appreciation to *David Asatiani, President of the Georgian Bar Association*, for his visionary leadership and unwavering commitment to advancing the legal profession in the era of digital transformation.

We also acknowledge the exceptional coordination efforts of *Tamta Devdariani, GBA Project Coordinator*, whose dedication ensured seamless collaboration between all stakeholders throughout the research and drafting process.

We are equally grateful to the legal practitioners and law firm representatives who participated in consultations and focus groups, generously sharing their insights, concerns, and innovative ideas. Their contributions provided the practical foundation for many of the recommendations outlined in this report – *Nika Sanadiradze, Sally Bezhashvili, Levan Makharashvili*.

Finally, we extend a note of appreciation to the Academician *Archil Prangishvili* for is invaluable insights and overall guidance. Also, *Isaac Asimov*, who's famous "Three Laws of Robotics" and broader body of considerations continue to inspire critical thought on the ethical dimensions of technology, a theme that resonates throughout this report.

** ChatGPT 4o model was used for translation, proof-reading and formatting means in this report.*

Table of Content

Acronyms and Abbreviations

ABA	American Bar Association
AI	Artificial Intelligence
AGI	Artificial General Intelligence
ANI	Artificial Narrow Intelligence
ASI	Artificial Superintelligence
CaC	Contract as Code (as programmatic code)
CPD	Continuing Professional Development
DSL	Domain-Specific Language
EU AI Act	European Union Artificial Intelligence Act
GBA	Georgian Bar Association
GDPR	General Data Protection Regulation
LLM	Large Language Model
M2M	Machine-to-Machine Integration
MCP	Model Context Protocol
ML	Machine Learning
RaC	Regulation as Code (as programmatic code)
RAG	Retrieval-Augmented Generation
SLA	Service Level Agreement
UPL	Unauthorized Practice of Law

Foreword: GBA

I am honored to present this report, which addresses one of the most critical questions facing our profession today: how to embrace the opportunities offered by digital transformation and Artificial Intelligence while preserving the ethical, procedural, and human foundations of justice.

This document is the result of close collaboration between GBA, Nastavia, and a wide range of practitioners whose contributions have been invaluable. It brings clarity to areas that are often misunderstood, from the mechanics of AI systems and their inherent risks to the structural reforms required for sustainable adoption. Equally important, the report charts a practical roadmap for modernizing our legal ecosystem—introducing digital tools to improve access to justice, developing AI governance standards, and strengthening professional capacity through education and training.

The insights and recommendations outlined here will guide our efforts in shaping a legal environment that is transparent, resilient, and aligned with international standards, while remaining deeply rooted in the principles of fairness and accountability. GBA is committed to maintaining the momentum of this process. We will work actively with state authorities, academic institutions, and international partners to ensure that these measures are implemented effectively and that our profession remains at the forefront of innovation without compromising its core values.

This report is not the conclusion of a discussion; it is the beginning of a long-term strategic journey for Georgia's legal profession. Together, we will ensure that technological progress strengthens, not erodes, the rule of law.

David Asatiani
President, Georgian Bar Association

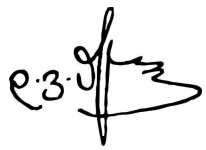

Foreword: Nastavia

I am both proud and deeply encouraged by the work presented in this report. What began as a thoughtful exchange of ideas with the respected partner, quickly evolved into a groundbreaking initiative that addresses one of the most pressing challenges of our time: how to harness the transformative power of digital technologies and Artificial Intelligence while safeguarding the principles that define our legal and ethical systems.

This report reflects more than research, it embodies collaboration, foresight, and shared commitment. We worked hand in hand with GBA leadership, legal practitioners, and institutional stakeholders to identify not only the extraordinary opportunities that digital transformation can unlock but also the critical risks that must be managed to maintain trust and integrity in the justice system. Together, we have mapped a path that is both ambitious and pragmatic: from embracing AI-driven innovation and streamlining judicial processes to building robust governance frameworks, investing in Georgian-language AI resources, and preparing future lawyers for a rapidly evolving digital landscape.

I would like to extend my heartfelt gratitude to the Georgian Bar Association for its vision and leadership, to our dedicated project team for their relentless effort, and to every practitioner and expert who contributed their insights. Your voices have shaped a report that is practical, future-ready, and uniquely tailored to Georgia's legal ecosystem.

At Nastavia, we believe that technology must always serve humanity, not replace its core values. This report stands as a testament to that belief, a blueprint for a legal profession that is innovative, resilient, and ethically grounded. It is our hope that the recommendations herein will not only guide action but inspire confidence in a future where digital progress and the rule of law advance hand in hand.

Teimuraz Murgulia
Founder, Nastavia

Executive Summary

This project began as an insightful dialogue between the President of the Georgian Bar Association and Nastavia's Head of Digital Transformation on the profound opportunities and risks presented by digitalization in the era of Artificial Intelligence (AI). What started as a focused exchange of ideas quickly evolved into a recognition that these issues demand structured research and a strategic vision. The conversations revealed a complex reality: while AI promises efficiency and innovation, it also challenges the legal profession with ethical dilemmas, regulatory gaps, and systemic risks that cannot be ignored.

Driven by this understanding, GBA and Nastavia launched a joint initiative to explore these dimensions in depth. The project combined expert analysis with practitioner perspectives, engaging GBA staff, individual lawyers, and representatives of law firms in structured discussions and focus groups. Their active participation provided invaluable insights into both the practical pain points of today's legal workflows and the strategic concerns for tomorrow's AI-driven ecosystem. These findings shaped the foundation of this report, which not only diagnoses challenges but also charts a path toward solutions that GBA, in partnership with state institutions and international allies, can transform into impactful and future-proof actions.

The first theme explored in the report is *digital transformation as an enabler of unlimited opportunities*. Lawyers expressed a strong demand for modernization across multiple fronts: remote and hybrid participation in court proceedings, AI-powered transcription of hearings, digital exchange of documents with public agencies, and real-time tracking of document status within workflow management systems. Requests extended to queue management tools, electronic signatures, common calendars for coordinating hearing dates, and online verification of lawyer credentials. These are not merely conveniences, they represent a fundamental shift toward efficiency, transparency, and accessibility in justice delivery. Nastavia's recommendations build on these needs, advocating for mandatory machine-to-machine integration for government services, methodologies for assessing the trustworthiness of AI service providers, and exploration of advanced technologies such as smart contracts and Contract as Code (CaC) ("Code" means programmatic code in this case, not the "codex" from the legal terminology). Together, these steps lay the groundwork for a modernized legal ecosystem, aligned with global best practices while tailored to Georgia's specific context.

The report then turns to *modern AI risks and challenges*, emphasizing that innovation cannot be separated from responsibility. This section demystifies core AI concepts, such as the distinction between ANI, AGI, and

ASI, and explains why limitations like hallucinations, data cutoffs, and embedding errors matter in legal settings. It further explores systemic concerns including algorithmic bias, lack of explainability in AI outputs, and vulnerabilities in data privacy when using third-party or cloud-based systems. Issues of regulatory ambiguity and client mistrust emerge as pressing challenges, alongside the heightened risk for smaller-language jurisdictions like Georgia, where limited digital corpora reduce model accuracy and reliability.

Building on this foundation, the report presents *sector-specific AI-driven risks from GBA's perspective*, focusing on challenges uniquely amplified within the legal profession. These include threats to attorney–client privilege when interacting with external AI systems, the potential erosion of legal craftsmanship as generative drafting tools replace traditional research methods, and unauthorized practice of law enabled by AI platforms. GBA's consultations also revealed concern over competitive disparities: large firms equipped with advanced AI capabilities may gain disproportionate advantages over smaller practices, while junior lawyers face shrinking opportunities to acquire core skills, an issue already evident in other sectors that rushed into automation. These trends demand proactive strategies to maintain both professional integrity and equitable access to digital tools.

Recognizing these realities, the final part of the report outlines a system framework and capacity-building roadmap for sustainable implementation. This includes ethical and operational standards for AI use, GBA-led certification mechanisms for AI tools, and compliance checklists for law firms. It also emphasizes the development of a national legal data governance framework, investments in Georgian-language AI resources, and phased adoption of machine-readable compliance through RaC, starting with Contract as Code as a practical entry point. Equally important is capacity building: equipping lawyers with AI literacy through continuing professional development, integrating technology and ethics into law school curricula, and preparing the judiciary and regulators to evaluate AI-driven processes. The roadmap, detailed in Appendix A, follows six structured phases, from awareness to continuous oversight, ensuring a balanced, risk-based approach to adoption.

As the aim of the team was to produce actionable, reference document, this report concludes with recommendations, enabling GBA and its partners to move from analysis to implementation. These measures aim not only to mitigate risks but to position Georgia's legal profession as a forward-thinking leader in the region, one that embraces innovation while safeguarding the values of fairness, accountability, and trust at the heart of the rule of law.

A. Digital Transformation: Unlimited Opportunities

Digital transformation represents the strategic use of technology to streamline institutional processes and fundamentally redefine delivery of services. Unlike some misconceptions, it is not limited to digitizing existing workflows; rather, it involves rethinking how entire systems operate in a connected, data-driven world. At its core, digital transformation seeks to reduce inefficiencies and eliminate redundant manual steps. In case of the legal sector, it is about seamless ecosystem where lawyers, courts, clients, and public institutions can interact securely and effectively. Respectively, this means moving beyond paper-based processes and siloed systems toward integral platforms that enable real-time collaboration, transparency, and automation, ultimately reinforcing the rule of law while adapting to the demands of a rapidly evolving digital society.

The core success of the project was the exploration of unlimited space of opportunities for digital transformation in the legal sector. Feedback from Georgian legal practitioners and law firms underscores a strong demand for modernization, both in the courts and across related public services, complemented by Nastavia's strategic vision for integrating cutting-edge technologies. Taking together, these priorities reveal a strong focus toward a connected, transparent, and intelligent legal ecosystem.

Digitalization of Court Processes

Modern court systems increasingly adopt remote and hybrid participation models, enabling lawyers to attend hearings virtually when physical presence is impractical. This reduces travel burdens, improves scheduling flexibility, and ensures continuity of proceedings during disruptions. Complementing remote access, lawyers highlighted the need for audio recordings of court sessions, extended with AI-based transcription tools (in parallel with existing text protocols). These solutions allow practitioners to maintain accurate case records and improve preparation for appeals or compliance checks.

Further take on scheduling would be a shared scheduling mechanism, providing the possibility to see the Judge's agenda (public part of it) and allow the practitioner to adjust and request reservation of free slots.

A court decision online registry further enhances transparency, granting immediate access to judgments and related documentation. Combined with online information portals for extradition cases, such systems reduce the time and cost of information retrieval. Together, these measures significantly streamline court interaction and align Georgia with global trends in digital justice.

Streamlined Workflow and Document Management

Effective case handling depends on efficient document flow management systems. Lawyers requested platforms that allow them to track the status of submitted documents in real time, reducing the uncertainty and delays that currently affect procedural steps. Introducing digital exchange of documents with law enforcement agencies and prosecutors would eliminate repetitive paper-based processes, accelerate responses, and strengthen compliance with procedural deadlines.

To reinforce this ecosystem, electronic signature capabilities should be mandated as a legal alternative for all official exchanges, ensuring security and legal validity while reducing physical dependencies. Further extension of the common calendar system which may span across courts, public agencies, and legal practitioners would allow real-time synchronization of hearing dates and notifications, minimizing scheduling conflicts.

Improving Access to Public and Institutional Services

Lawyers frequently encounter administrative delays that could be alleviated by queue management systems and online appointment booking for visits to courts and public agencies. Extending this principle, enabling remote client meetings in correctional facilities - where confidentiality is not compromised - would reduce logistical barriers and improve the efficiency of attorney-client interactions. Similarly, secure online meetings with prosecutors would support pre-trial negotiations and streamline procedural coordination.

An essential addition to this transformation is online verification of lawyer credentials, reducing the risk of fraud and ensuring trust in digital interactions across courts and institutions.

Collaboration and Knowledge Sharing

Digital transformation also offers an opportunity to strengthen professional collaboration. Lawyers expressed interest in an online platform for peer discussion and knowledge exchange, which could include thematic forums, legal updates, and practice tips. Such a platform would foster collective problem-solving and serve as a resource hub for smaller practices with limited access to specialized expertise.

Leveraging AI and Advanced Technologies

Among the most transformative elements is the integration of AI into legal practice. Lawyers indicated growing interest in exploring AI capabilities for research, drafting, and predictive analytics, recognizing its potential to accelerate workflows. However, AI adoption must be supported by strong ethical and operational safeguards, an issue addressed later in this report.

B. Modern AI Risks and Challenges

Foundation for the Regulatory-Aware AI

Today, the roadmap to effective AI adoption begins with a critical first step - recognizing that, beyond a small circle of tech professionals, most people, including those in the legal field, have only a limited understanding of how AI truly works and how outputs are generated by Large Language Models (LLMs) or Machine Learning (ML) systems. While legal practitioners are not expected to master the technical details of neural networks or training processes, the absence of a basic conceptual understanding creates two major risks: the improper use of AI tools and vulnerability to misleading claims from so-called "AI experts." In an environment where AI-powered solutions are marketed aggressively, these risks are further amplified by the hype surrounding emerging technologies.

This section highlights the essential considerations that every law firm and legal practitioner should understand when adopting AI or managing client cases. AI is not merely a tool; it redefines the operational landscape across nearly every profession. In some areas, it may lead to complete disruption, while in most others, it represents a profound paradigm shift that demands careful adaptation rather than passive acceptance.

AI Categories (ANI, AGI, ASI)

Artificial Intelligence is often discussed as if it were a single concept, yet the tech world distinguishes between three major categories: Artificial Narrow Intelligence (ANI), Artificial General Intelligence (AGI), and Artificial Superintelligence (ASI). ANI refers to systems designed for specific tasks within a limited domain, such as legal document review or language translation. These systems excel at narrowly defined objectives but lack the capacity to transfer knowledge beyond their programmed scope, making them powerful yet fundamentally limited. AGI represents the theoretical next stage, a form of intelligence capable of performing any intellectual task a human can, with reasoning, learning, and problem-solving abilities across diverse domains. Achieving AGI would mark a profound shift, as such systems could adapt to entirely new contexts without retraining, a capability far beyond what is possible today. No existing system has reached AGI, though research continues toward this goal with growing intensity. ASI, by contrast, refers to a hypothetical intelligence surpassing human cognitive capabilities in virtually every area, including scientific discovery, strategic planning, and creative reasoning. While AGI and ASI remain speculative, they frame the ethical and regulatory debates surrounding AI, highlighting the vast gap between today's task-specific tools and the transformative systems that may emerge in the future.

Core AI concepts

First, the way AI systems do their job, significantly differs from the general information systems and expectations linked with them. In the traditional tools, if the user looks for the "civil code", it returns an exact match if it is present in the database and no result if it is not. So if you don't enter the exact term or citation (or at least, the proper pattern), you may get no results. AI-powered systems work differently. They convert words and concepts into mathematical “vectors” (embeddings) and search for the closest meaning, not just exact matches, and the definition of "closest" does not always correspond to the human's expectations.

This means AI can return relevant information even if the wording is different. Clearly, this leads to the great flexibility and usability of the system but also, because AI relies on similarity rather than certainty, it can “hallucinate” - generate answers that sound convincing but are factually incorrect or unsupported. And people who are used to using AI systems based on the "exact match" paradigm, may easily be misled by such results. Most models have a possibility to configure the level of similarity through settings like temperature (how creative or varied responses are). Lower temperature gives more predictable answers; higher temperature introduces variability - useful for brainstorming but risky for formal legal advice.

Another important limitation is the training data cutoff date - best AI models are trained on thousands of special processors, during hundreds of days, so neither of them are naturally "up to date". The best and largest models as of 06/2025, usually have cutoff dates somewhere between late 2023 and early 2024.

So, anything published after that, such as the latest court decisions or newly adopted regulations, is missing from the model’s knowledge. To address this, modern systems use techniques like Retrieval-Augmented Generation (RAG) and emerging Model Context Protocol (MCP). In simple terms, these act like a real-time legal assistant with an up-to-date library: instead of relying only on what the AI “remembers,” they fetch the most current documents from trusted sources and feed them into the AI for analysis. For example, if a lawyer needs the latest amendments to specific law, the AI can retrieve the official text first, then summarize or compare it, ensuring that advice is based on the actual law, not outdated knowledge.

Another important point is that not all languages are equal for AI systems. Large Language Models (LLMs) perform best in languages with abundant, high-quality training data such as English, French, or Spanish and even other types of GenAI tools, like voice or image generation, are better suited for such languages.

For smaller languages like Georgian, two major limitations arise. First, data scarcity: there are far fewer digitized legal texts, court decisions, and professional writings in Georgian compared to dominant languages, making it harder for the model to learn accurate syntax, terminology, and context. Second, embedding challenges: vector representations (used for semantic search and reasoning) are often optimized for widely spoken languages, which can lead to lower precision and subtle misinterpretations in Georgian. This means that AI-generated outputs in Georgian may lack nuance, omit critical legal terms, or even misstate provisions - posing additional risks for lawyers relying on these tools without rigorous review.

Data Quality and Algorithmic Bias

AI systems are fundamentally dependent on the data they consume, and this dependence introduces significant risks when the underlying information is incomplete, inconsistent, or biased. Most modern AI models are trained on vast historical datasets, which inevitably contain imperfections and human biases. When such biases are embedded in the data, they can be reproduced or even amplified by the AI system. For instance, if historical hiring records reflect gender imbalances, an AI tool designed for recruitment assistance might replicate those patterns, leading to discriminatory recommendations. The more specific cases for AI ethics and bias are discussed in the follow-up chapter, but even though this simpler example lies outside the legal sector, it illustrates how biased AI outputs can easily escalate into compliance disputes, employment litigation, or regulatory scrutiny.

The challenge is not limited to bias alone. Data quality and governance play an equally critical role. In practice, information is often scattered across multiple sources, poorly structured, and inconsistently labeled. When an AI system processes unclassified or outdated material, the risk of inaccurate conclusions rises significantly. This is why strong governance practices such as data classification, cataloguing, and lineage tracking are indispensable. Classification ensures that sensitive or privileged information is properly identified and handled, cataloguing creates transparency about the origin and version of each source, and lineage provides clarity on which datasets influenced a given AI decision. Without these safeguards, it becomes extremely difficult to assess the reliability of outputs or to establish accountability when errors occur, issues that often surface in regulatory investigations and liability claims.

Compounding these challenges is the lack of explainability inherent in most large-scale AI systems. These models operate as complex "black boxes," generating outputs without offering insight into their internal

reasoning. This problem becomes even more acute when organizations rely on external AI services provided by third parties.

Proprietary models rarely disclose the logic behind their conclusions, leaving users uncertain whether a particular assessment, for example, an automated creditworthiness score, stems from sound financial indicators or from latent biases within the dataset. Such opacity can trigger legal and ethical concerns, especially when decisions impact individuals' rights or financial standing, and there is a growing trend for regulators becoming more demanding with organizations to demonstrate transparency in AI-driven processes.

AI Ethics, regulatory awareness and core risks

Nastavia's ongoing research on the concept of "Regulatory-Aware AI" emphasizes critical dimensions of AI ethics that will shape the trajectory of future development, particularly as technology moves from today's narrow applications toward the potential horizon of Artificial General Intelligence (AGI). Within this framework, the exchange of perspectives with the Georgian Bar Association has brought valuable insights, reinforcing the severity of risks and the need to embed ethical principles and regulatory considerations into AI systems from the earliest stages of their evolution.

The ethical behavior of current AI systems does not arise from a built-in understanding of moral principles but from patterns absorbed during training. Unlike the fictional framework - "Three Laws of Robotics", envisioned by Isaac Asimov, where a strict hierarchy of robotic laws ensured that higher-level rules always prevailed over lower-level ones, modern AI lacks such foundational algorithmic order. With the first law - "A robot may not injure a human being or, through inaction, allow a human being to come to harm" - placed above all others, guaranteeing that no subsequent rule could override it. In his novel *The Robots of Dawn*, Asimov, through the character Susan Calvin even remarked that a mind bound by these laws would make an ideal mayor because of its unerring obedience to rules and protection of human welfare. In contrast, real-world AI systems have no comparable hierarchical safeguards. Their behavior is driven by statistical correlations in training data, and even advanced alignment methods such as reinforcement learning from human feedback reflect subjective human judgments.

Consequently, what is deemed "good" or "acceptable" is highly contingent on the biases embedded within the datasets used for training. In case of imposing completely new regulations, this lag can turn into a significant risk. This fragility extends far beyond ethics into the realm of law and regulation. While large language models can process and summarize enormous volumes of legal material, they cannot fully grasp the

dynamic and hierarchical nature of legal systems. Regulations differ across jurisdictions, evolve, and often demand nuanced interpretation.

ANI lacks the cognitive capacity to apply proportionality tests, reconcile conflicts between overlapping norms, or assess the relative weight of different provisions within a legal hierarchy. Although these systems can retrieve relevant context, they cannot guarantee legally sound or consistent application. This gap becomes particularly critical in compliance-intensive domains, where overlooking a recent amendment or misinterpreting a rule could result in significant liability. One promising approach to mitigate this limitation is the adoption of Regulation as Code (RaC) – the practice of translating regulatory requirements into machine-readable formats that can be systematically processed by AI systems. Properly implemented, RaC can serve as a bridge between static legal texts and dynamic AI reasoning, enabling systems to access authoritative, up-to-date regulatory logic rather than relying solely on probabilistic interpretation.

Beyond these inherent limitations lies there is a more insidious risk: the deliberate or accidental embedding of a Triggered Disinformation Vector (or "Magnet Mines"). This refers to a latent behavioral mechanism within an AI system that activates only under narrowly specific conditions. Unlike general misinformation, which spreads broadly, this type of manipulation is selective and conditional. It may trigger based on a user's identity, geographic location, metadata, or interaction pattern, releasing false or strategically biased outputs only to a targeted audience. Imagine an AI-driven compliance assistant widely adopted by multinational corporations. Under normal conditions, it provides accurate guidance on tax obligations. However, for users from a particular jurisdiction, it subtly misstates a regulatory threshold, leading to underreporting. Such a scenario could create hidden vulnerabilities, expose companies to penalties, and even serve as a tool for geopolitical or economic sabotage. Because these triggers operate covertly, they are extremely difficult to detect, audit, or attribute, raising profound concerns about accountability and transparency.

Data Privacy and Confidentiality

Matters of data privacy and confidentiality are those most often voiced by legal professionals. While there is a clear understanding of the potential benefits of digital transformation, many participants in the focus groups expressed concerns about the practical aspects of managing sensitive information. These concerns typically fall into three critical areas: data storage, data processing, and data exchange, while the level of specific awareness varies significantly per area.

Storing digital data

With the increasing adoption of digital technologies, both in professional legal practice and general information systems, data security remains a pressing concern for many legal professionals - *how to keep records, in Paper or Digital form?* Today, most lawyers are generally aware of the broad spectrum of modern cybersecurity risks - including data breaches, unauthorized access, and system vulnerabilities, but many still lack the technical knowledge needed to envisage how effective safeguards work. For example, some lawyers responded they would never allow their clients' confidential data to be digitized in any form, opting instead to store all sensitive materials physically in a secure office safe or a bank deposit box.

Although storing paper documents in a safe may seem like a reliable method for protecting sensitive data, it is often based on a biased perception that equates physical storage with superior security, while viewing digital formats as inherently vulnerable. However, this comparison is often incomplete, as it overlooks critical aspects such as the continuous physical security of the office, the certified quality of the safe, the secure handling of keys and PINs, and the human factor in physical breaches.

A more balanced comparison begins by imagining both formats, a printed document and a USB drive, left unattended on a desk. In such a scenario, both are equally vulnerable. This illustrates that effective data security is less about the format and more about the protective measures applied. Just as one invests in reinforced locks and surveillance systems for physical files, securing digital records demands measures such as encrypted storage, firewalls, strong password policies, and access control mechanisms.

From a practitioner's perspective, it may seem more intuitive to trust the security of a physical office and the good safe, than to evaluate firewalls and encryption systems. While this assumption is understandable, it is not always well-justified, as both physical and digital safeguards require professional expertise and continuous oversight to withstand intrusion attempts, which becomes more sophisticated with time.

Local device vulnerability is considered among top cons for digitalization. Storing sensitive legal data on local devices - desktops, laptops, or external drives - introduces risks such as malware, ransomware, phishing, and insider threats. These attacks can lock or steal data, causing severe financial and reputational harm. However, implementing modern safeguards such as full-disk encryption, secure boot, multi-factor authentication, and advanced endpoint protection not only mitigates these risks but also delivers significant advantages. Properly secured devices provide lawyers and especially larger legal firms with immediate, controlled access to case files without relying on physical archives, while encrypted storage ensures that even if a device is stolen, the data remains protected.

It is easy to define who can see what and actually oversee how those rights were used. Combined with automated backups and remote-wipe capabilities, these measures create a level of resilience and continuity that physical storage cannot match. It is also important to note that social engineering threats, such as those exploiting data leaks, apply to both digital and paper-based records, yet managing these risks is far easier with robust digital rights management and activity monitoring tools.

When it comes to using *cloud-based storage*, which offers a significant level of convenience, redundancy and scalability, another side of the medal - critical risks like data residency, third-party security vulnerabilities, and uncontrolled replication across multiple servers shall be considered. Sensitive client information stored in cloud environments may be subject to foreign jurisdictions or regulatory frameworks that conflict with local data protection laws, creating potential compliance issues. However, larger cloud providers put enormous efforts to comply with recognized security standards, offering contractual guarantees on data residency, and enabling encryption for data both in transit and at rest. In general Implementing private or hybrid cloud models further strengthens control over storage locations. These measures not only reduce legal and ethical exposure but also provide clear advantages - including improved disaster recovery, remote accessibility, and the ability to implement advanced security features like end-to-end encryption and audit trails, which are often beyond the reach of traditional on-premise setups. Thus, perceiving cloud storage as simply putting files in some unknown land is a misconception - when properly planned and managed, it is more like installing your own secure safe, with custom code, inside a high-security datacenter, which is often far better protected than the local bank branch.

One of the most overlooked yet critical aspects of data security is the *Data Retention & Deletion Strategy*, or, simply put, how client information is retained and deleted when no longer needed. The risk lies in residual data stored in archives, system logs, or backup copies that can be unintentionally exposed during breaches or audits. Inadequate deletion practices may lead to compliance violations under data protection laws, including GDPR and Georgian regulations, which require strict adherence to retention periods. These challenges can be addressed through clear retention policies, automated deletion workflows, and backup solutions that support secure erasure protocols. Modern storage systems also offer version control and data lifecycle management features, which not only help maintain compliance but deliver operational advantages - reducing storage costs, streamlining document management, and minimizing the risk footprint over time. To conclude, proper deletion policies do not just prevent problems - they actively enhance efficiency and governance across the practice.

While data privacy and confidentiality remain a critical issue for legal professionals, modern tech development has overcome most traditional weaknesses. While risks still exist, they can be effectively managed through proper governance, technical measures, and informed practices. resulting in similar (if not more) safety than physical methods, while adding extra level of opportunities. As digital transformation is no longer a luxury, but a matter of competitive advantage, especially when it comes to AI - *own data becomes a key success factor to effectively utilize AI capabilities*.

Data processing

Compared to data storage, the risks associated with data processing, especially in non AI systems, often receive less attention, yet they are equally critical. This chapter highlights the challenges that arise when client information is handled through various software solutions, including both AI-driven and traditional systems. "Data at work" is no less prone to security risks than "Data at rest".

As previously outlined, the growing reliance on software and cloud-based services in legal practice introduces new *Software and Compliance Risks*. Tools for document or case management, client relationship management or even simple collaboration services may process sensitive client data on external servers, raising concerns about data breaches, unauthorized access, and uncertainty over where and how information is stored.

Additionally, non-compliance with regulations such as GDPR, Georgian law on "Personal Data Protection", or professional codes of ethics can lead to severe legal and reputational consequences. These risks, however, are manageable through the consistent security focus and structured oversight. Selecting vendors that adhere to internationally recognized standards such as ISO 27001 or SOC 2, coupled with clear contractual clauses on data residency, encryption, and processing limitations, provides a strong baseline of security and compliance. Implementing regular vendor audits and due diligence processes further reduces exposure.

When handled properly, using specialized legal software offers clear advantages - from streamlining document workflows and reducing manual errors to enabling secure, traceable collaboration across teams and jurisdictions. Compliance-focused platforms often include built-in features such as access/audit logs, automated retention policies, and real-time monitoring, giving firms more control and transparency than traditional manual processes ever could. In short, with careful vendor selection and active governance, third-party software can become a powerful enabler of efficiency and compliance rather than a source of risk.

Another critical consideration is the choice of software itself. Many non-Ai systems are also a “black box” which accepts inputs and delivers outputs,

while its internal logic and data handling practices remain largely unknown. As a result, relying on improperly vetted software, regardless of its stated security levels, poses a serious risk of information leaks and breaches.

Data exchange

Secure data exchange is one of the most critical aspects of maintaining client confidentiality in the digital era. Legal professionals routinely share sensitive information with clients, courts, prosecutors, law-enforcement agencies and third-party providers, making transmission channels a potential point of failure. While this offers speed and convenience, it also opens the door to a range of security risks if not properly managed.

Despite the emergence of machine based (M2M) direct data exchange protocols, email remains the preferred medium for legal correspondence, yet *default email protocols are vulnerable by design* and many company owned email servers are not properly configured. Messages sent without encryption can be intercepted, altered, or misdirected, while file-sharing links from unverified platforms often lack adequate access controls.

Data exchange issues become even more important, when data streams or packets cross state borders and go international. These jurisdictional complexities can create uncertainty about which standards apply and who bears liability in the event of a breach. For example, messaging apps without end-to-end encryption expose conversations to potential breaches because of challenges with transnational regulations.

While data transmission risks are common sense today and even people far from the tech sector are aware to "look for the lock" while working with the their browsers (which is an established indicator that SSL/TLS transport security protocol is enabled), there are other threats, bearing much higher levels of sophistication.

Man-in-the-Middle Attacks (MITM) are sophisticated techniques designed to effectively mislead users about the identity of the counterpart they are communicating with. These attacks often involve creating malicious proxy servers that can intercept and modify data in transit or even redirect traffic entirely, giving users the illusion that they are interacting with a legitimate resource when, in fact, they are not. In some cases, attackers deploy fully mocked user interfaces to make the deception even more convincing.

While the simple practice of “look for the lock” (checking for HTTPS and secure certificates) remains one of the most effective defenses against such attacks, it is often inconsistently applied. Unlike financial services platforms, where users are trained to verify security indicators by default, in case of many non-financial ones, including eGovernment services, this is not applied in a similarly consistent way, while the importance of verification may be equally high or even greater for some cases.

Regulatory and Legal Ambiguity

The deployment of AI across professional services is advancing far faster than the development of clear regulatory frameworks, resulting in legal grey zones that raise profound questions about accountability, liability, and professional standards. When does an AI system deliver inaccurate or biased outcomes, whether in law, medicine, accounting, or consulting, who is responsible - the provider that built the model, the professional who relied on it, or the client who acted on the advice? Who is the owner of GenAI outputs? What about free vs paid plans? Current legislation provides no definitive answers, leaving practitioners and organizations exposed to uncertainty and potential litigation.

Beyond these ambiguities, AI introduces a broader structural shift in how intellectual professions compete. Historically, fields such as law, medicine, and financial advisory relied primarily on human expertise, judgment, and reputation. While institutional resources offered advantages, an exceptionally skilled individual could often rival large firms through deep knowledge and personal competence. AI is reshaping this equation. Advanced AI systems are becoming the equivalent of production assets in industrial sectors, powerful tools that dramatically improve efficiency, accuracy, and scale. Just as no individual can match the output of a modern factory without access to comparable infrastructure, professionals without AI capabilities will find it increasingly difficult to compete with organizations that have heavily invested in these technologies.

In the short term, this shift may appear beneficial for companies to take advantage over solo entrepreneurs in Georgia, where AI can help local firms improve service delivery and reduce costs. However, as AI systems evolve, adapt to smaller languages, and integrate regulatory logic, barriers to entry will favor well-capitalized global players. International firms equipped with advanced AI capabilities will find it easier to enter local markets, offering highly automated and competitively priced services. These dynamic risks widening the gap between large institutions and independent professionals, not only in law but across all knowledge-based sectors.

Without regulatory safeguards and equitable access to AI tools, the competitive landscape may consolidate around a small number of dominant actors, raising concerns about market concentration, client choice, and the long-term sustainability of smaller practices.

Client mistrust

As AI becomes integrated into professional services, including law, it introduces a new dimension of trust dynamics between clients and practitioners. At the current stage of AI development, two opposing attitudes are emerging, each carrying significant risks. On one end of the

spectrum, some clients view AI tools as inherently unsafe or unreliable, insisting that their cases or projects be handled without any AI involvement. While understandable from a risk-averse perspective, such demands can weaken a company's ability to deliver services efficiently and competitively, particularly when peers leverage AI for research, analysis, and decision support.

On the opposite end, certain clients develop an overinflated perception of AI’s capabilities, expecting near-perfect accuracy and faster results without appreciating the technology’s limitations. When outcomes fail to meet these unrealistic expectations, blame often falls on the service provider, even if the error lies in the inherent unpredictability of AI systems or the client’s misunderstanding of its role. Both extremes underscore the urgent need for clarity in how AI is deployed within client engagements. In sectors like legal services, this clarity should be reinforced through strong ethical frameworks, transparent communication, and explicit contractual provisions governing AI use. Establishing such protective layers is critical not only for maintaining client confidence but also for mitigating disputes, safeguarding professional accountability and aligning AI adoption with regulatory and ethical standards.

C. Sector specific AI-Driven Risks and Challenges: GBA Take

With proper understanding of core AI concepts and systemic challenges, we can turn to the sector specific matters, requiring direct attention from the professional regulator. While AI tools offer undeniable efficiencies, they introduce a spectrum of challenges that are uniquely amplified in the legal field. Law operates within strict ethical frameworks and confidentiality requirements, where even minor errors can have severe consequences.

For Georgia, as outlined, these dynamics are compounded by additional factors. Risks extend far beyond technical failures: they include potential breaches of attorney–client privilege, over-reliance on generative systems that lack contextual reasoning, and the erosion of foundational skills for young lawyers. The growing disparity between large firms with resources to adopt advanced technologies and smaller practices struggling to keep pace further threatens market balance and access to justice.

The objective is clear: to provide GBA with actionable insights that safeguard professional integrity while enabling responsible innovation - one that positions the Bar not as a passive observer, but as a leader in shaping the future of the legal profession in the AI era.

Hindering Entry into the Legal Profession

As AI becomes embedded in legal practice, one of the most underappreciated yet critical risks is its potential impact on entry pathways for new professionals. The Georgian Bar Association must examine lessons from other sectors, particularly technology, where rapid adoption of AI tools has significantly altered workforce dynamics. In many technology companies, the pursuit of "fast" AI-driven efficiency has led to the elimination of junior roles traditionally responsible for repetitive or routine work. While this appeared to optimize operations in the short term, it created an unintended consequence: removing the very foundation upon which professional development rests. Junior staff, once entrusted with basic research, drafting, and process tasks, are often the talent pool from which future experts emerge.

This trend appears in the legal profession as well. A notable article by William Josten, published by the Thomson Reuters Institute in 2025 under the title "Why the Changing Composition of Law Firms May Pose a Problem for GCs," highlights a direct connection between the growing adoption of generative AI tools and a projected decline in the proportion of junior staff within law firms. The analysis warns that while these technologies promise efficiency, they risk disrupting the traditional talent pipeline, creating long-term challenges for developing future legal expertise.

The experience in the tech sector shows that even the most advanced AI systems, including today's Artificial Narrow Intelligence (ANI), remain tools rather than full replacements for human judgment. Companies that aggressively downsized junior teams based on unrealistic expectations of AI autonomy soon encountered operational bottlenecks.

For instance, in 2023, several large software companies, including those in Silicon Valley, cut entry-level coding positions after integrating AI-assisted development tools such as GitHub Copilot. Within a year, these firms faced delays and rising costs because they lacked trained developers capable of handling complex projects, debugging AI-generated code, and managing architecture-level decisions, tasks that AI tools could not reliably perform. Some companies were forced to reverse these cuts, rehiring or retraining staff at significant expense. Ex. There is a case with Klarna, online payment services provider, which admitted mistakes with firing nearly 700 employees from customer service and marketing units.

Report from the recognized organizational design and workforce planning platform, Orgvue, named "Human-first, machine enhanced: From optimism to pragmatism in AI-driven workforce transformation" has an important finding: "55% of businesses admit wrong decisions in making employees redundant when bringing AI into the workforce".

The legal profession faces an even higher level of complexity and risk. Unlike code, which can be validated through testing, legal reasoning involves interpretive nuances, contextual weighting of norms, and ethical considerations that AI cannot replicate. If the law firm rapidly eliminates opportunities for early-career lawyers to engage in foundational work such as case research, drafting, due diligence - the profession may experience a hollowing out of talent. This could result in a future shortage of experienced practitioners capable of exercising independent judgment, undermining both the quality of legal services and public trust.

Ehe challenge is amplified in Georgia by the limitations of local-language AI training data. As current systems perform less accurately with smaller languages, than in English or other high-resource languages, there is an increasing the likelihood that AI-generated outputs will require scrutiny by human lawyers. This makes it even more dangerous to reduce junior capacity prematurely, as the human review layer is indispensable for now.

The message for GBA and its members is straightforward: meaningful and managed AI adoption is not optional, it is inevitable. However, it must be structured in a way that preserves pathways for professional growth, maintains quality standards, and ensures resilience against the overestimation of AI capabilities. Developing guidance for firms on workforce planning, training programs for AI-augmented workflows, and ethical standards for balanced adoption should form a core component of GBA's strategic agenda.

Fragmentation of Legal Advice

The growing integration of generative AI into legal workflows introduced new possibilities for speed and efficiency, yet it also poses a significant risk: the fragmentation of legal advice and the erosion of legal craftsmanship. Unlike traditional drafting processes, which rely on structured reasoning and a comprehensive understanding of legal principles, AI-driven outputs often prioritize textual similarity over contextual precision. This can lead to fragmented legal opinions, inconsistencies in interpretation, and in some cases, outright fabrication of authority.

In a highly publicized incident, a New York attorney faced sanctions after submitting a legal brief that cited multiple nonexistent cases generated by ChatGPT. The lawyer admitted to failing to verify the authenticity of these citations, resulting in a $5,000 fine and a mandate to inform all affected clients about the misconduct. Similarly, a Texas lawyer was fined $2,000 for filing a document containing fictitious citations produced by an AI tool. During the discussion of these cases, including the ways to prevent such behavior in the local legal practice, the project team paid special attention to the fact that the court additionally required the attorney to complete a course on the ethical use of generative AI in legal settings and noted the creation of the similar setup in Georgia as an another possible priority for the Georgian Bar Association. These cases, widely reported by the New York State Bar Association (nysba.org), illustrate a profound challenge: AI is not a substitute for professional judgment, and the failure to treat it as an assistive tool rather than an autonomous solution can lead to reputational damage and professional discipline.

The risks extend beyond individual cases. If practitioners rely on AI outputs without rigorous human oversight, the legal profession faces a long-term threat of diminished analytical depth and reduced standards of argumentation. Junior lawyers traditionally develop their skills through extensive engagement with statutes, law, and client context - processes that require time and intellectual effort. Overuse of generative drafting tools risks bypassing these foundational steps, resulting in a generation of practitioners who can operate AI platforms but lack interpretive judgment. This shift creates fragmented advice, where documents may appear polished but lack the coherence and legal integrity expected by clients.

These concerns are reflected in the American Bar Association’s Formal Opinion 512, issued in July 2024, which provides the first formal framework for the ethical use of generative AI in legal practice. The opinion underscores that the Model Rules of Professional Conduct remain fully applicable in the context of emerging technologies, and it calls for lawyers to exercise heightened vigilance when integrating AI tools into client representation.

The opinion draws attention to core obligations under the Model Rules:

> *Competence (Rule 1.1) requires lawyers to understand the benefits and risks of generative AI and to verify the accuracy of its outputs. AI "hallucinations"—plausible but fabricated text—are now a well-documented phenomenon, making unverified reliance on these tools incompatible with professional competence.*
>
> *Confidentiality (Rule 1.6) remains a cornerstone of legal ethics. Because generative AI systems often require inputting sensitive client information, lawyers must ensure that the tools they use provide robust data security. If confidentiality cannot be guaranteed—such as in cases where AI providers use input for further model training—lawyers must refrain from using the tool without the client's informed consent.*
>
> *Communication (Rule 1.4) obliges lawyers to keep clients informed of significant factors affecting representation. The decision to employ generative AI in preparing documents or research is material enough to warrant discussion, especially if it could affect cost, timing, or legal strategy. Transparency not only upholds ethical duties but also mitigates client mistrust and unrealistic expectations regarding AI capabilities.*

The fragmentation of legal advice is not an inevitable outcome of AI integration, but avoiding it requires deliberate strategies. Firms should embed AI tools within structured workflows that mandate human review and verification, rather than replacing these steps entirely. Training programs must emphasize not only technical literacy but also the enduring value of legal reasoning, ensuring that AI serves as an amplifier of professional competence rather than a substitute for it.

Unreliable software and services

Analysis of the software tools and online services to manage day-to-day operations in the legal sector reveals a large variety of types. These include document drafting platforms, legal research databases, case management and workflow systems, client communication tools, e-discovery platforms, and cloud-based storage solutions. In addition, new categories of AI-driven legal assistants and generative document automation services are entering the market, often marketed as productivity boosters capable of drafting contracts, summarizing case law, or predicting litigation outcomes.

While these tools promise efficiency, they also introduce critical risks when their reliability and security are uncertain. The integration of AI into these solutions significantly amplifies existing vulnerabilities. Document drafting software using generative AI can produce clauses that sound legally sound but are inaccurate, incomplete, or even contradictory to jurisdictional norms. For instance, a clause generated for an English law agreement might unintentionally appear in a Georgian contract template if the system fails to apply contextual restrictions correctly. Similarly, AI-powered legal research tools may present "hallucinated" case citations, as demonstrated by high-profile international incidents where lawyers faced sanctions for submitting briefs with fictitious references produced by ChatGPT-like systems.

Case management systems that incorporate AI for automated deadline tracking or compliance alerts may misinterpret procedural rules, particularly in smaller jurisdictions like Georgia, where localized legal data is scarce. A missed court deadline caused by an inaccurate AI-generated notification could have severe consequences for both lawyers and their clients. Furthermore, cloud-based storage and document exchange platforms introduce confidentiality risks when integrated with third-party AI services that process sensitive client data without transparent data usage policies. This raises concerns about attorney, client privilege, data sovereignty, and GDPR compliance.

The rise of AI-driven predictive analytics and contract review services also brings the risk of over-reliance on algorithmic recommendations. Lawyers may accept risk scores or compliance flags without questioning the underlying logic, that often operates as a "black box," with no guarantee of explainability or jurisdiction-specific accuracy. This not only jeopardizes case strategy but also complicates liability: who is accountable when an AI-generated recommendation leads to a negative outcome, the lawyer, the firm, or the software vendor?

At present, there is no formal mechanism to vet these tools before adoption, leaving individual practitioners and firms to make decisions without standardized benchmarks for security, accuracy, or ethical compliance. In an environment where some providers exaggerate capabilities while concealing risks, the absence of oversight creates fertile ground for malpractice, reputational harm, and erosion of client trust.

To address this challenge, the Georgian Bar Association can establish a dedicated AI Ethics and Technology Task Force, which may review key risk factors, product proposals, ethical and technical initiatives and provide respective recommendations. Such a structure will not only safeguard the profession against unreliable software, but also set a precedent for proactive governance, reinforcing GBA's role as the guardian of professional integrity in the digital era.

D. System Framework and Capacity Building

The resulting judgement of the project can be outlined as a following summary - addressing the risks associated with AI in the legal sector requires more than individual best practices, it calls for a systemic approach that combines regulatory clarity, ethical standards, and professional preparedness. The rapid pace of AI adoption, coupled with legal ambiguity and uneven access to advanced tools, demands a structured framework to ensure that innovation strengthens, rather than undermines, the integrity of the legal profession. This involves creating clear rules for AI use, embedding transparency and accountability into professional conduct, and aligning national policies with international benchmarks such as the EU AI Act, GDPR and notable rulings from other BAR associations.

Capacity building is a cornerstone of this transformation. Lawyers, regulators, and judges must be equipped with the knowledge and tools necessary to manage AI responsibly, while law schools need to integrate technology and ethics into their curricula. At the same time, investments in data governance and local language resources are critical for ensuring that AI systems perform accurately in the Georgian context. By combining these efforts with phased adoption strategies, certification mechanisms, and collaborative knowledge-sharing platforms, Georgia can establish a resilient, regulatory-aware AI ecosystem that enhances competitiveness without compromising ethical and legal standards.

Establishing a Regulatory-Aware AI Framework

Recognizing the importance of embedding regulatory and ethical (also as set of rules) foundations in the AI, Nastavia started working on the "Regulatory-aware AI" voluntary standard at the early emergence stage of machine learning and language models. Avalanche-like development of the killing tech, based on AI, triggered by the Russian invasion in Ukraine, while understandable, further rises concerns about the ability to keep those developments under control where necessary. The standard, being aimed for global application is similarly usable on a national scale. Thus, the joint effort with the GBA to push forward structured framework would position Georgia as a regional leader in responsible AI governance but also provide the legal sector with clear principles for integrating AI tools without compromising professional ethics or regulatory compliance

Regulatory-Aware AI ecosystem means embedding compliance and ethical safeguards directly into the design and deployment of AI systems, ensuring that outputs respect legal hierarchies, privacy norms, and professional accountability standards. This approach shifts the narrative from reactive regulation to proactive governance, where AI solutions are built with transparency, auditability, and explainability as default features.

For Georgia, this is especially critical given the rapid modernization of its legal environment and its trajectory toward EU integration. Alignment with international frameworks such as the EU AI Act and GDPR must serve as a cornerstone of this initiative. These regulations establish principles for risk-based AI classification, human oversight, and data protection, all of which are vital for legal applications where confidentiality and fairness are paramount. Incorporating these standards into a Georgian context, while considering local language challenges and market dynamics, will require adaptive strategies, including consultation with legal professionals, regulators, and technology providers. By doing so, Georgia can establish a practical model for regulatory-aware AI governance that strengthens trust, enhances competitiveness, and mitigates systemic risks.

Sector-specific Ethical and Operational Standards for AI Use

The absence of well-defined ethical and operational standards for AI in legal practice creates serious risks that have already appeared globally, as highlighted in earlier chapters discussing fabricated citations and breaches of attorney-client privilege. Thus, regulations and respective enforcement mechanisms help law firms to avoid inconsistent practices, client mistrust, and exposure to liability. To address this gap, the Georgian Bar Association should take the lead in creating GBA-endorsed AI ethics guidelines, setting uniform expectations for law firms and practitioners across the country.

These guidelines must define mandatory principles: transparency in disclosing AI usage to clients, accountability for outputs regardless of tool involvement, informed client consent where sensitive data is shared with external systems, and a requirement for human oversight at every stage of AI-assisted work. Building on these pillars, the introduction of AI audit trails should become standard practice, enabling verifiability of how AI-generated content was produced and reviewed. This mechanism is critical not only for resolving disputes, but for maintaining professional integrity in an era where accountability can become obscured by automation.

In addition, standards should address vendor risk, requiring that firms adopt only those AI tools that meet minimum compliance benchmarks for security, confidentiality, and explainability. These requirements should extend to clear disclosure of data usage policies, ensuring that client information is not repurposed for model training without explicit consent, and to auditability mechanisms that allow independent verification of compliance. By operationalizing these principles, GBA can create a framework where innovation enhances, rather than undermines, the rule of law and client trust. Inaction, by contrast, risks allowing fragmented practices to proliferate, weakening the profession’s ability to uphold ethical obligations in an increasingly AI-driven environment.

Institutional and Professional Capacity Building

The successful integration of AI in legal practice is not only about tech, it requires comprehensive capacity building across the profession and its institutions. Without long term and deliberate investment in knowledge, skills, and oversight structures, the risks outlined in earlier chapters, ethical lapses, client mistrust, and erosion of legal craftsmanship, will only deepen.

The first priority is *AI literacy for legal professionals*. Lawyers have to develop operational familiarity with AI tools, in par with an understanding of how these systems work, their limitations, and their inherent risks. This includes knowledge of key concepts such as generative AI, hallucinations, data governance, and confidentiality, as well as an awareness of ethical obligations and regulatory frameworks. Integrating these topics into *mandatory Continuing Professional Development (CPD) programs* will ensure that all practitioners reach a minimum standard of competency for responsible AI adoption.

At the same time, the profession must address the challenge of *skill transition pathways for junior lawyers*. The automation of routine tasks reduces opportunities for early-career practitioners to build foundational skills through traditional research and drafting. Eliminating these opportunities risks creating a generation of lawyers who can operate AI platforms but lack independent analytical judgment. GBA should promote alternative pathways, such as supervised workflows where juniors validate AI outputs, and AI-assisted research labs where they learn both the benefits and limitations of emerging technologies. These strategies preserve the intellectual precision of legal training while preparing young professionals for an AI-augmented future.

Equally critical is *judiciary and regulator preparedness*. Courts and oversight bodies should be informed about AI-generated submissions, while ensuring compliance with established ethical norms, and addressing issues of admissibility, confidentiality, and transparency. Specialized training for judges and regulators is essential, alongside practical tools for auditing AI systems and investigating misconduct. Aligning these capabilities with top international frameworks, will help Georgia maintain consistency with its broader legal harmonization goals.

Finally, *embedding capacity building into the legal ecosystem* requires collaboration among GBA, law schools, and technology partners. Academic curricula must integrate AI ethics and legal technology, while initiatives such as legal tech clinics, simulation exercises, and innovation incubators can give students and practitioners practical exposure. These efforts reinforce the central principle that, even in an AI-driven environment, human judgment remains indispensable.

Data Governance and Local Language Development

The reliability of AI systems is only as strong as the quality of the data on which they operate. In legal practice this becomes a critical mission. For the AI user, it means prioritizing the development of a data governance framework that ensures structured, accurate, and secure handling of legal information. Such a framework must define standards for data storage, data processing, safeguarding confidentiality throughout AI workflows; and data exchange, enabling interoperability without compromising privacy. Without these measures, the adoption of AI risks amplifying inaccuracies and breaching confidentiality, leading to both ethical violations and regulatory exposure.

A critical component of this effort is the creation of high-quality, structured legal datasets for the Georgian language. It is an upcoming trend for modern and highly flexible businesses, to look for fast and efficient contracting mechanisms to avoid disputes at all. The success of smart contracts is just emerging, but this would further improve. In the absence of proper mechanisms in the local legislation, more businesses would tend to deal on external platforms or even in external jurisdiction with efficient dispute resolution and enforcement. To close this gap, Georgia needs to launch Georgian Language AI Development Initiatives, leveraging public–private partnerships to build comprehensive legal corpora, encourage anonymized data sharing from law firms, and incentivize technology providers to invest in Georgian-specific model fine-tuning. These steps will not only improve the reliability of AI outputs but also ensure equitable access to the benefits of AI across the legal profession.

Beyond data curation, Regulation as Code (RaC), so called Domain Specific Language (DSL) offers a transformative approach to embedding compliance in AI systems. By converting complex regulatory frameworks into machine-readable formats, RaC enables AI tools to apply authoritative legal logic rather than relying solely on probabilistic interpretation. This shift could substantially reduce compliance risks and make regulatory processes more efficient. However, the effectiveness of RaC depends on its implementation strategy. Global RaC standards offer advantages of consistency and interoperability, allowing multinational systems to operate under harmonized principles. Yet they often lack the granularity needed for language-specific and jurisdiction-specific nuances, which can lead to gaps in interpretation or misalignment with local law. Conversely, local RaC implementations provide precise alignment with national legislation and linguistic contexts but may limit cross-border compatibility and require continuous updates as laws evolve.

A practical starting point for this evolution is Contract as Code (CaC) DSL, which may use keywords to formalize contractual obligations in a way that machines can interpret and validate.

DSL Keyword	Semantic Meaning
PERMIT	Specifies an action or obligation that is allowed under the contract.
PROHIBIT	Defines an action that is explicitly forbidden by the agreement.
OBLIGATE	Creates a binding requirement for one party to perform an action.
CONDITION	Establishes a rule that must be true for an obligation or permission to apply.
DEADLINE	Sets a specific date or time limit for fulfilling an obligation.
FULFILLED	Declares condition fulfilled.
FAILED	Declares condition failed.
PENALTY	Defines a financial or procedural consequence triggered by a breach or delay.
TRIGGER	Specifies an event that activates an obligation, right, or penalty.
TERMINATE	Provides the condition under which the contract or a clause ends automatically.
NOTIFY	Represents the requirement to send formal notice under certain conditions.
DISPUTE	Establishes rules for conflict resolution, including arbitration or escalation steps.

The table shows possible example keywords and their semantic meaning for CaC use. This approach can serve as the foundation for future RaC frameworks, allowing legal professionals to experiment with rule-based automation in a controlled scope before scaling to broader regulatory contexts. CaC can streamline routine transactions by automatically verifying compliance with key contractual provisions, monitoring deadlines, and triggering alerts for breaches or performance obligations. Common use cases include automated loan agreements, where interest adjustments and penalty clauses can execute dynamically, service-level agreements (SLAs) that monitor uptime guarantees and apply remedies in real time, and procurement contracts, ensuring that delivery terms and payment milestones are enforced without manual intervention. These use-cases demonstrate how structured legal logic, when encoded as rules, can significantly reduce errors, accelerate transactions, and build confidence in automated legal processes.

For Georgia, the optimal solution lies in a hybrid approach: adopting global RaC principles for structural consistency while building localized layers that reflect the Georgian legal system and language. By beginning with CaC and gradually expanding toward full RaC implementation, stakeholders can manage complexity, ensure language-specific accuracy, and build institutional readiness for advanced regulatory automation. This strategy ensures that AI systems can deliver compliance-ready outputs tailored to domestic requirements without isolating the jurisdiction from international legal and technological standards. By leading such initiatives, GBA can position Georgia at the forefront of regulatory innovation in the region, turning language and jurisdictional specificity from a barrier into a strategic advantage.

Certification and Compliance Mechanisms

The ongoing trend for AI tools to become integral to legal practice, the need for a structured certification and compliance system is critical. Without standardized vetting processes, law firms risk adopting tools that lack adequate safeguards for security, confidentiality, and ethical use. To address this, the Georgian Bar Association may establish a GBA-driven AI tool review program specifically tailored to legal practice. Such certification would validate that a tool meets minimum compliance benchmarks for data privacy, explainability, human oversight, and adherence to relevant regulatory frameworks, including GDPR and emerging EU AI Act principles.

Alongside such review, law firms should follow a mandatory compliance checklist before integrating any AI system into their workflows. This checklist should cover core questions such as: Does the tool ensure client data confidentiality? Are AI outputs subject to human review? Is there a documented mechanism for error reporting and accountability? Is the vendor transparent about data usage policies and model training practices? Incorporating this step into procurement decisions will prevent ad hoc, inconsistent adoption that lead to ethical breaches and regulatory penalties.

Vendor transparency is another cornerstone of this approach. AI providers must disclose details regarding data retention, processing, and model governance, as well as whether user inputs are stored or reused for training. Requiring such disclosures as a condition for certification will not only enhance trust but also create a culture of accountability among technology vendors. By institutionalizing these mechanisms, GBA can help build a legal ecosystem where innovation coexists with strong compliance, protecting both practitioners and their clients from the systemic risks associated with opaque AI adoption.

Sustainable Implementation Roadmap

Achieving a regulatory-aware, ethically grounded AI ecosystem in the legal sector requires a phased, systematic approach rather than isolated initiatives. The roadmap outlined in Appendix A: High-Level Roadmap provides the foundation for this strategy, identifying six interdependent phases that guide the profession from awareness to long-term governance.

Phase 1 focuses on awareness and foundational capacity building, ensuring that legal professionals develop a baseline understanding of AI fundamentals, associated risks, and ethical obligations. This includes AI literacy programs, GBA-issued guidelines for responsible adoption, and early-stage consultations with law firms to assess current practices.

Phase 2 introduces policy and governance structures, creating the normative framework for ethical AI use. This involves drafting a dedicated AI Ethics and Professional Responsibility Code, establishing certification schemes for AI tools, and piloting Regulation as Code (RaC) initiatives in collaboration with regulatory bodies.

Phase 3 addresses data governance and language-specific challenges, recognizing that AI systems cannot operate reliably without structured, high-quality legal data. A national legal data governance framework, combined with efforts to build a Georgian-language legal corpus, will ensure that AI outputs reflect both accuracy and jurisdictional context.

Phase 4 expands into professional training and academic integration, embedding AI competencies across the legal ecosystem. Mandatory Continuing Professional Development (CPD) modules, law school curricula updates, and AI-assisted research clinics will ensure that both current and future practitioners acquire the skills necessary for effective oversight and ethical use of AI technologies.

Phase 5 centers on regulatory alignment and risk-based adoption, harmonizing Georgia's standards with international frameworks such as the EU AI Act and GDPR. This phase includes defining high-risk AI applications in legal practice and introducing a compliance certification system for law firms and technology vendors.

Phase 6 ensures sustainability through continuous monitoring and collaboration, including the creation of a dedicated AI Oversight Committee, engagement in international bar associations' AI task forces, and publication of an Annual AI Risk and Innovation Report.

The roadmap is designed as a living instrument. While Appendix A provides the initial draft structure, further development will require stakeholder consultations, resource planning, and continuous refinement to align with both global best practices and the specific needs of Georgia's legal community.

Recommendations

The insights and strategies outlined in this report highlight both the immense opportunities and the complex risks associated with AI adoption and digital transformation in the legal sector. To ensure these findings translate into measurable progress, the Georgian Bar Association can take a proactive role in sustaining momentum, institutionalizing best practices, and driving a forward-looking AI agenda. The following recommendations aim to create a structured foundation for implementing the proposed initiatives, safeguarding professional integrity, and positioning Georgia's legal profession as a leader in ethical, technology-enabled innovation:

R1. Accelerate Digital Transformation Initiatives

GBA can strengthen collaboration with government institutions, courts, and technology providers to accelerate implementation of modernization priorities identified by practitioners, such as machine exchange, use of digital documents and electronic signatures, remote collaboration and remote presence instruments, in parallel with emerging AI technologies, while ensuring alignment with professional obligations.

R2. Institutionalize AI Ethics and Oversight

GBA can introduce AI ethics as a mandatory component of Continuing Professional Development (CPD) and establish a dedicated "AI Ethics and Technology Task Force". This body can monitor emerging risks, review legal tech solutions, and issue guidance and recommendations to uphold ethical and operational standards.

R3. Prevent Barriers to Entry for New Professionals

GBA can design policies and structured programs, such as supervised AI workflows, AI-assisted research labs, and mentorship initiatives, to preserve learning pathways for junior lawyers. These efforts will help maintain professional sustainability and ensure that analytical skills remain central to legal practice.

R4. Promote Regulatory-Aware AI policies

GBA may find it necessary to support the development and adoption of Regulatory-Aware AI principles both nationally and through international dialogue. This includes participating in standard-setting initiatives, advocating for Regulation as Code (RaC) approaches adapted for Georgian law, and collaborating on policies that embed compliance and transparency into AI systems by design.

Appendix A. High Level Roadmap

Phase	Objective	Key Actions	Responsible Stakeholders	Timeline
Phase 1: Awareness & Foundation	Build basic understanding of AI risks, benefits, and ethical obligations within the legal community.	- Launch GBA-led AI Literacy Campaign (seminars, webinars, guides). - Publish **GBA Guidelines on Ethical AI Use** (aligned with ABA Model Rules & EU AI Act principles). - Initiate consultations with law firms to assess current AI usage.	GBA, Bar Committees, Legal Tech Experts	
Phase 2: Policy & Governance Framework	Create systemic standards for safe and transparent AI use.	- Draft **AI Ethics & Professional Responsibility Code** for legal practice. - Establish **AI Tool Certification Scheme** for law firms. - Advocate for **Regulation as Code (RaC)** pilot projects with Ministry of Justice.	GBA, Ministry of Justice, National Data Protection Agency	
Phase 3: Data Governance & Language Support	Develop high-quality legal data infrastructure and support Georgian language AI capability.	- Create a **Legal Data Governance Framework** (classification, cataloguing, lineage). - Launch **Georgian Legal Corpus Project** in partnership with academia & tech sector. - Set rules for **secure data exchange** to preserve attorney–client privilege.	GBA, Universities, Tech Firms, Judiciary	
Phase 4: Professional Training & Academic Integration	Equip lawyers and future practitioners with AI competency.	- Develop **Mandatory CPD Module on AI and Ethics** for all licensed lawyers. - Integrate **AI Law & Technology Courses** into law school curricula. - Create **AI Legal Clinics** to give students practical experience under supervision.	GBA, Law Schools, Judicial Training Centres	
Phase 5: Regulatory Alignment & Risk-Based Adoption	Ensure compliance with international standards and enable structured adoption.	- Align national standards with **EU AI Act** and GDPR. - Define **high-risk AI use cases in legal practice** and mandate additional safeguards. - Introduce **AI Audit & Compliance Certification** for law firms.	GBA, Parliament, Data Protection Agency	
Phase 6: Continuous Monitoring & International Collaboration	Maintain regulatory awareness and foster knowledge exchange.	- Establish **AI Oversight Committee** within GBA. - Join **international bar associations' AI task forces**. - Launch an **Annual GBA AI Risk & Innovation Report** to track trends.	GBA, International Partners, Tech Regulators	