



**REGULATING AI-SUPPORTED LEGAL GUIDANCE
IN DIGITAL LEGAL EDUCATION PLATFORMS: A
RISK-BASED FRAMEWORK FOR YOUTH LEGAL
LITERACY**

Aminjon Qalandarov

*PhD in Law, Acting Associate Professor,
Tashkent State University of Law, Tashkent, Uzbekistan*

E-mail: aminjonqalandarov9@gmail.com

ARTICLE INFO:

Received: 01.10.2026

Revised: 02.10.2026

Accepted: 03.10.2026

KEY WORDS:

*artificial intelligence;
legal advice; legal
education; legal
literacy; youth; AI
governance; risk-based
regulation; Uzbekistan*

ABSTRACT

Artificial intelligence (AI) can expand access to legal knowledge by making legal information more searchable, conversational, and responsive to user needs. However, as digital legal education platforms begin to provide fact-specific explanations, personalized guidance, or recommendations, the boundary between legal education, legal information, and individualized legal advice becomes increasingly difficult to maintain. This challenge is particularly important for children and young users, who may be more vulnerable to legal error, over-reliance, opaque personalization, and unnecessary processing of personal data. This article develops a risk-based framework for AI-supported legal guidance in digital legal education platforms. Using doctrinal, comparative, and design-oriented methods, it examines the EU Artificial Intelligence Act, the Council of Europe Framework Convention on Artificial Intelligence, UNESCO and UNICEF guidance, the OECD AI Principles, professional ethics guidance, empirical research on legal hallucinations, and relevant Uzbek legislation and policy. The analysis also draws on the scientific-practical concept of the “Huquq maktabi 2.0” project. The article proposes four functional risk tiers: educational explanation, legal navigation, personalized legal guidance, and high-stakes legal advice. It argues that safeguards should increase in proportion to personalization,



legal consequence, and user vulnerability. The proposed model combines authoritative-source grounding, transparency, human oversight, data minimization, age-appropriate design, referral mechanisms, and redress. Rather than relying on a binary distinction between “information” and “advice,” the article advocates a functional approach capable of supporting legal literacy while reducing the risks of unsafe automated legal guidance.

1. Introduction

Digital legal education is moving from static repositories of legislation and explanatory materials toward interactive systems that can answer questions, recommend content, assess knowledge and personalize learning pathways. Generative AI accelerates this transition because it allows users to communicate with legal content in ordinary language rather than through legal terminology or formal search logic. In principle, this can lower cognitive and informational barriers to law, especially for young users who may not know which legal source, institution or procedure is relevant to a problem. The same interface, however, can create a new regulatory difficulty: a system designed to educate may begin to function as an adviser even when its institutional mandate, technical architecture and safeguards were not designed for individualized legal decision-making.

The difficulty is not semantic. Legal information generally describes the law, legal education aims to build understanding and capability, while legal advice applies law to a person’s specific circumstances and can influence consequential decisions. Conversational AI blurs these categories because the interaction is iterative, personalized and responsive. A user can move within seconds from asking what an employment contract is to asking whether a particular dismissal is lawful, whether a deadline has expired, or what document should be filed tomorrow. If the system responds with an apparently authoritative answer, the user may reasonably rely on it even where the system was intended only as an educational tool. This risk is amplified for children and young people because age, experience and digital literacy can affect how strongly conversational outputs are trusted and how well uncertainty is understood (UNICEF, 2025; UNESCO, 2023).

Current AI governance increasingly adopts risk-sensitive principles rather than treating all AI uses identically. The OECD AI Principles emphasize human-centered, robust, transparent and accountable AI (OECD, 2024). UNESCO’s Recommendation on the Ethics of Artificial Intelligence places human rights, dignity, fairness, transparency and human oversight at the center of responsible AI governance (UNESCO, 2022). The Council of Europe Framework Convention establishes a treaty framework under which Parties are to implement risk and



impact management responsive to potential effects on human rights, democracy and the rule of law (Framework Convention on Artificial Intelligence, 2024). The EU Artificial Intelligence Act adds mandatory transparency requirements for systems that interact directly with natural persons and establishes a broader architecture in which regulatory intensity is related to context and risk (Artificial Intelligence Act, 2024). These instruments do not create a single ready-made rule for a youth legal education chatbot, but together they support a functional approach in which safeguards should increase as personalization, legal consequence and vulnerability increase.

Legal-domain evidence reinforces the need for caution. Large language models can perform impressively on legal examinations and structured tasks (Katz et al., 2024), but performance on benchmark tasks does not eliminate factual and legal unreliability. Dahl et al. (2024) found high rates of legal hallucination in public-facing general-purpose models and showed that such systems may accept false premises rather than correct them. Evaluations of specialized legal research systems also demonstrate that retrieval augmentation reduces but does not eliminate erroneous or unsupported outputs (Magesh et al., 2025). These findings matter most where users lack the expertise to identify an incorrect answer. Professional ethics guidance likewise emphasizes competence, confidentiality, verification and human responsibility when generative AI is used in legal practice (American Bar Association, 2024).

Uzbekistan provides a particularly relevant setting for this inquiry. The national Strategy for the Development of Artificial Intelligence Technologies until 2030 calls for the formation of a regulatory framework for AI based on international experience and for the development of population skills and institutional capacity (Resolution No. PQ-358, 2024). In parallel, national policy on continuous legal education for minors expressly supports modern visual tools, videos, computer games and mobile applications that enable schoolchildren to acquire legal knowledge independently (Resolution No. 532, 2024). The “Huquq maktabi 2.0” project builds directly on this policy environment and envisages a personalized, interactive, AI-supported legal learning platform with a chatbot and, later, a voice assistant. Crucially, its scientific-practical concept frames AI as an auxiliary mechanism for legal navigation, explanation and content recommendation rather than an autonomous source of independent legal conclusions.

This article asks a practical regulatory question: how should AI-based legal guidance within a digital legal education platform be designed and governed so that it improves legal literacy without quietly becoming unsafe automated legal advice? The analysis proceeds from the view that the relevant boundary is functional rather than terminological. It therefore distinguishes educational explanation, legal navigation, personalized guidance and high-stakes advice, and then connects those functions to proportionate safeguards. The resulting model is applied to the institutional setting of a youth-oriented platform such as “Huquq maktabi 2.0,” where legal education, personalization and conversational AI intersect in a single user environment.



The research problem sits at the intersection of three bodies of scholarship that are often discussed separately: the reliability and hallucination risks of legal AI, the boundary between legal information and professional legal advice, and the pedagogy of AI-supported legal education for younger users. Existing studies illuminate each of these areas, but they offer less guidance on a public legal-education platform in which the same conversational interface can move, during one interaction, from a general explanation to fact-sensitive guidance. The contribution of this article is to connect interaction type, legal consequence and user vulnerability to a graduated set of safeguards that can be translated into platform design and governance.

2. Literature and Regulatory Context

The emerging literature on generative AI and law contains a productive tension. On one side, advanced language models demonstrate substantial capability on legal reasoning and examination tasks. Katz et al. (2024) reported that GPT-4 achieved a score above the passing threshold across Uniform Bar Examination components. Such performance suggests that language models can support legal learning, issue spotting, explanation and document-oriented tasks. On the other side, empirical studies reveal that fluency is not reliability. Dahl et al. (2024) documented pervasive legal hallucinations and found that models may answer confidently even when the legal premise of a question is incorrect. Magesh et al. (2025) further showed that specialized AI legal research tools using retrieval-augmented generation remain capable of producing erroneous information. The policy implication is not that legal AI should be excluded from education, but that user interface fluency must never be treated as evidence of legal accuracy.

Legal education scholarship increasingly treats AI competence as part of professional and learning outcomes. Kimbrough (2025) argues for integrating technology-driven lawyering competencies into legal education rather than treating AI as an external tool. For youth legal literacy, however, the educational objective is different from professional legal training. The goal is not to teach users to delegate legal judgment to an AI system; it is to help them recognize legal issues, locate authoritative sources, understand basic rights and duties, identify appropriate institutions, and know when professional assistance is required. This distinction is central to responsible platform design. An educational chatbot should therefore be evaluated not merely by whether it gives a plausible answer, but by whether the interaction improves the user's legal capability and directs the user toward verifiable sources and safe next steps.

International soft-law instruments support this capability-oriented approach. UNESCO's 2023 Guidance for Generative AI in Education and Research recommends a human-centered and age-appropriate approach, strong privacy safeguards, ethical validation and careful pedagogical design (UNESCO, 2023). UNICEF's updated guidance on AI and children likewise emphasizes safety, privacy, non-discrimination, transparency, explainability, accountability, inclusion and the best interests and development of children (UNICEF, 2025). These principles



are particularly important in legal education because a chatbot may receive sensitive descriptions of family conflict, violence, school discipline, employment, migration, health, finances or other personal problems. A platform that merely adds a generic disclaimer while collecting such information would not adequately address the structural risk.

The EU AI Act provides additional legal architecture. Article 50 requires providers of AI systems intended to interact directly with natural persons to ensure that users are informed that they are interacting with AI, unless that fact is obvious in context. Article 5 prohibits certain manipulative or deceptive AI practices and specifically addresses exploitation of vulnerabilities linked to age where the relevant conditions are met. Annex III treats certain AI systems in education and vocational training as high-risk when they perform specified functions such as determining access, admission or assignment, although an ordinary legal literacy chatbot is not automatically high-risk merely because it is educational (Artificial Intelligence Act, 2024). The important analytical point is that the Act differentiates risk by function, consequence and context rather than by the generic label placed on the system.

The Council of Europe Framework Convention complements this approach by linking AI governance to human rights, procedural safeguards, transparency and iterative risk assessment (Framework Convention on Artificial Intelligence, 2024). It is especially useful for legal-information systems because the quality of legal guidance affects practical access to rights and remedies. A person who is misdirected about a limitation period, complaint channel, eligibility criterion or procedural requirement may lose an opportunity even if no formal decision was made by the AI system. In legal education, therefore, 'risk' should include not only discrimination or privacy harm but also legal-pathway harm: the possibility that inaccurate or overconfident guidance causes a user to choose the wrong legal route, delay action or misunderstand the status of an authoritative source.

Professional ethics materials, although directed to lawyers rather than educational platforms, provide a useful benchmark for the highest-risk end of the spectrum. ABA Formal Opinion 512 emphasizes that lawyers using generative AI remain responsible for competence, confidentiality, communication, supervision, candor and verification (American Bar Association, 2024). This supports an important design principle: when a digital platform crosses from general education into individualized, consequential legal advice, it should not treat model output as self-authenticating. Human professional responsibility, source verification and clear allocation of accountability become increasingly necessary.

A related strand of scholarship addresses the regulatory boundary between legal information and the practice of law. Avery, Abril, and del Riego (2024) show that generative AI destabilizes traditional unauthorized-practice-of-law categories because software can now perform functions that resemble individualized legal services without fitting conventional professional structures. Ryan and Hardie (2024), testing generative AI on common legal problems, likewise identify material risks when non-lawyers rely on machine-generated legal information and advice. More



recent analyses argue that legal regulation should neither treat all automated assistance as prohibited practice nor assume that consumer-facing AI can safely provide unrestricted advice; instead, the regulatory focus should turn to function, consequence, competence, supervision and access to justice (Comoglio, 2024; Glesner Fines, 2025). This literature supports the present article's decision to avoid a purely label-based distinction between 'information' and 'advice.'

Uzbek law and policy provide a complementary domestic framework. The Law on Personal Data defines personal data broadly and requires legality, accuracy, confidentiality, protection and respect for constitutional rights in personal-data processing (Law on Personal Data, 2019). For a personalized legal education platform, these principles support data minimization and purpose limitation: the system should not collect names, identifiers or sensitive factual histories merely because conversational interfaces make such collection easy. The 2030 AI Strategy calls for a robust national legal basis for AI and a unified system for regulating relations arising from broad AI deployment (Resolution No. PQ-358, 2024). The policy on continuous legal education for minors, meanwhile, demonstrates an explicit state interest in digital and mobile legal-learning tools (Resolution No. 532, 2024). Together, these instruments create both an opportunity and a responsibility: AI-supported legal literacy can be institutionally encouraged, but its safeguards should be defined before high-stakes use becomes normalized.

Uzbekistan's Law on Advocacy is also relevant to the functional boundary examined here. Article 5 identifies, among the forms of legal assistance provided by advocates, consultations and explanations on legal issues, oral and written information on legislation, and the drafting of legal documents (Law on Advocacy, 1996). This provision does not itself resolve whether or when an AI system could constitute unauthorized professional practice, and the present article does not claim that it does. It does, however, demonstrate that individualized legal consultation and document-oriented assistance belong to a legally recognized professional domain. That makes it important for a publicly supported educational platform to define *ex ante* when automated interaction must remain educational, when it may provide constrained guidance, and when it should refer the user to a qualified professional or competent authority.

3. Methodology

This study uses a combined doctrinal, comparative and design-oriented methodology. The doctrinal component identifies legal and normative requirements relevant to AI-mediated legal education: transparency, human oversight, risk assessment, privacy, age-appropriate design, professional responsibility and reliability. Primary normative materials include the EU Artificial Intelligence Act, the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, Uzbekistan's AI Strategy to 2030, Uzbekistan's Law on Personal Data, and national measures on continuous legal education for minors. International policy instruments from UNESCO, UNICEF and the OECD are treated as persuasive governance standards rather than binding domestic law.



The comparative component does not attempt to transplant any one foreign regulatory regime wholesale. Instead, it identifies convergent principles across systems and asks which of them are functionally relevant to a legal education platform operating in an emerging AI regulatory environment. The study therefore distinguishes between (a) binding rules that directly regulate an identified context, (b) professional standards that illuminate the responsibilities attached to legal advice, and (c) soft-law principles that can guide system design before specific national legislation exists.

The design-oriented component converts these principles into a risk-tier framework. The unit of analysis is not the AI model as a whole but the user-facing function performed in a particular interaction. This is necessary because the same technical model may perform low-risk educational explanation in one turn and high-risk individualized guidance in the next. The framework therefore classifies interactions by five cumulative variables: degree of personalization, legal consequence, time sensitivity, user vulnerability, and factual dependency. Safeguards are intensified as these variables increase.

The article also uses the “Huquq maktabi 2.0” scientific-practical concept as an applied case study. The project concept envisions legal lessons, tests, quizzes, mini-cases, segmented content, personalization, an AI-supported legal guidance module, a chatbot and a voice assistant. It explicitly states that AI should function as an auxiliary mechanism for legal navigation, explanation and content recommendation rather than as an autonomous provider of independent legal conclusions. This institutional choice provides a strong starting point for a risk-based architecture because it places legal capability and safe referral above automated adjudication or autonomous advice.

The proposed framework is normative and architectural rather than an empirical effectiveness study. It does not claim that any particular technical implementation currently achieves the reliability described here. Nor does the analysis treat the proposed safeguards as additional legal or contractual obligations already imposed on the existing “Huquq maktabi” application or on the project team. Except where a requirement follows independently from applicable law or the approved project documents, the safeguards discussed below are presented as design options and regulatory recommendations for future development. Empirical work would still be needed to test user comprehension, error rates, source-grounding performance, differential effects across age groups and the effectiveness of referral mechanisms.

4. From Legal Education to Automated Legal Advice: A Risk-Based Model

A binary rule that separates ‘legal information’ from ‘legal advice’ is too crude for conversational AI. In static media, the difference is often visible: a website publishes general information, while a lawyer advises a client after obtaining facts. In a chatbot, the transition is gradual. The user supplies facts, the system asks follow-up questions, the explanation becomes tailored, and the output may recommend a course of action. A platform may continue calling the



interaction ‘education’ even though its practical effect is advisory. Regulation should therefore focus on function rather than label.

At the lowest-risk end is educational explanation. The system explains legal concepts, institutions and procedures using general examples. It may define a contract, explain the difference between civil and criminal liability, or describe the general purpose of a complaint mechanism. This function supports legal literacy and ordinarily requires no personal data. The principal risks are outdated law, oversimplification and failure to distinguish binding law from educational commentary. These can be controlled through source grounding, update protocols, links to authoritative materials and concise uncertainty language.

The second function is legal navigation. Here the system helps a user identify the relevant legal topic, public institution, official form, source or procedural pathway. Navigation can be highly valuable because citizens often do not know where to begin, but a wrong route can also waste time or contribute to a missed opportunity. In the model proposed here, a navigational response would ideally identify the relevant jurisdiction and date, refer to an authoritative source, flag the possible relevance of deadlines or eligibility conditions, and offer an official or human referral where the issue is time-sensitive or uncertain.

The third function is personalized legal guidance. The system uses facts supplied by the user to explain how general legal rules may relate to the user’s situation. At this point the interaction is no longer purely pedagogical, because a fact-sensitive inference may influence what the user does next. The proposed model therefore associates this function with stronger safeguards, including collection of only the facts needed for the interaction, a clear indication that the output is not a final professional opinion, source-grounded explanation, visible uncertainty and access to referral where the issue concerns possible loss of rights, sanctions, safety, minors, significant financial consequences or complex procedural choices.

The fourth function is high-stakes or outcome-determinative legal advice. Examples include recommending whether a claim should be filed, assessing the legality of a specific act in a way likely to be relied upon immediately, preparing a final filing without review, or advising in circumstances involving child safety, liberty, status, significant property interests or access to a remedy. For a platform whose primary purpose is legal education, the framework proposed here favors human review or referral at this point. AI may still organize information, identify sources, summarize options or prepare questions, but the consequential judgment would not be treated as an autonomous educational output.

This continuum also shows why a disclaimer, by itself, cannot resolve the boundary problem. Calling an output ‘educational’ does not necessarily change the practical effect of an interaction in which detailed facts are collected and a specific course of action is recommended. A more credible approach is to combine disclosure with functional limits: as the interaction becomes more individualized and consequential, the system can narrow the response, rely more heavily



on verified sources, communicate uncertainty more clearly, or redirect the user to an appropriate human or official channel.

4.1. Operationalizing the Four-Tier Model

The framework proposed here organizes interactions by risk rather than treating every AI-generated sentence in the same way. Its logic is cumulative: as personalization, factual dependence and legal consequence increase, additional safeguards become appropriate. The aim is proportionality. Basic educational explanation should remain usable and accessible, while more individualized guidance calls for stronger source, privacy and referral protections.

Table 1. Risk tiers for AI-supported legal guidance

Risk tier	Functional description	Primary risk	Recommended safeguards within the proposed framework
Tier 1: Educational explanation	General concepts, rights, duties, institutions; no user-specific legal conclusion	Outdated or oversimplified information	Authoritative-source links; update date; clear distinction between law and examples; plain language
Tier 2: Legal navigation	Directs users to relevant law, institution, procedure, form or official channel	Misdirection, missed deadlines, wrong jurisdiction	Jurisdiction/date checks; official-source grounding; deadline alerts; referral option; interaction disclosure
Tier 3: Personalized legal guidance	Applies general rules to user-provided facts without issuing a final professional opinion	Hallucination, over-reliance, privacy harm, incorrect factual inference	Data minimization; source-grounded responses; uncertainty labels; red-flag detection; human escalation; logging and quality review
Tier 4: High-stakes legal advice	Outcome-determinative recommendation or	Loss of rights/remedies, safety or status	No autonomous final advice; qualified human



	action with material legal consequences	harm, professional- responsibility failure	review or authoritative referral; enhanced consent; auditability; incident response and redress
--	--	--	---

Tier 1 is the natural starting point for a youth legal education platform. In this mode, the chatbot functions as a pedagogical interface to verified legal content rather than as an unrestricted source of answers. A prudent implementation would generate or check responses against legislation, official explanations and institutionally approved learning material, while indicating the date or version of a source where that matters. Because the purpose is educational, the interface can explain a rule, illustrate it with examples and direct the user back to the authoritative material rather than presenting a bare conclusion.

Tier 2 adds a navigational function. Where a user asks which authority deals with a matter, what form may be relevant or how a procedure begins, the quality of routing becomes as important as the quality of explanation. For an Uzbekistan-focused platform, a sensible design choice would be to prioritize official national legal sources and competent institutions, and to avoid presenting precise deadlines or procedural requirements unless they are grounded in a current authoritative source. Where the route is uncertain, the safer response is to indicate that uncertainty and point the user toward an official channel.

Tier 3 raises different concerns because personalization affects both privacy and reliance. The model therefore favors data minimization, separation of short-term conversational context from persistent user profiling, and source-grounded explanations that distinguish the user's stated facts from the applicable legal rule. Personalization for learning purposes can rely on information such as age group, language preference or course progress without turning a user's private legal problem into a lasting profile attribute. Where the system is uncertain, that uncertainty is most useful when it appears next to the relevant proposition rather than only in general terms of service.

At Tier 4, the proposed model shifts from automated guidance toward human review or authoritative referral. The AI can remain useful by summarizing the issue, identifying relevant documents, explaining general options or helping the user formulate questions for a lawyer or public body. What changes is the status of the final judgment: the platform would not present an unreviewed, outcome-determinative conclusion as though no further professional assessment were needed. This is consistent with the professional-responsibility logic reflected in ABA Formal Opinion 512, under which generative AI may assist legal work but does not displace human responsibility for accuracy, confidentiality and judgment (American Bar Association, 2024).



4.2. Proportionate Safeguards for Youth-Oriented Legal AI

The risk tiers become meaningful only when they are connected to concrete safeguards. This article argues that, within the proposed framework, responsible youth-oriented legal AI should rest on four core safeguards: reliable and current legal sources; transparent and age-appropriate interaction; privacy protection combined with referral for higher-risk matters; and continuing evaluation and accountability. This is a normative claim about sound system design. It does not mean that every element described below is already a contractual obligation of the existing “Huquq maktabi” application or the grant team.

Reliable legal guidance should begin with verifiable legal sources. Within the proposed framework, a platform should draw primarily from a curated corpus of current legislation, official guidance and institutionally approved educational materials, and should make the source of a legal proposition visible where this is reasonably possible. Version information is also essential because a legally correct answer can become inaccurate after amendment. Where a relevant source has changed or cannot be verified, the system should narrow or qualify its response rather than fill the gap from model memory. This requirement follows from the documented persistence of hallucination in both general-purpose and specialized legal AI systems (Dahl et al., 2024; Magesh et al., 2025).

Transparency and age-appropriate design should be treated as substantive safeguards rather than as optional interface features. Article 50 of the EU AI Act illustrates the broader principle that users should understand when they are interacting with AI (Artificial Intelligence Act, 2024), while UNICEF (2025) emphasizes safety, privacy, fairness, transparency and children’s best interests. In a youth legal-literacy setting, the system should therefore explain what it can and cannot do, avoid manipulative engagement practices, and discourage unnecessary disclosure of sensitive information. The form and depth of these explanations should be adapted to age and context instead of being confined to a lengthy privacy notice.

Privacy and referral become especially important once the system moves from general explanation to personalized guidance. Uzbekistan’s Law on Personal Data requires legality, accuracy, confidentiality and protection in data processing (Law on Personal Data, 2019). The proposed framework therefore distinguishes learning-profile information, such as language preference, completed modules or quiz performance, from detailed narratives about family, medical, financial or disciplinary problems. Where a conversation raises serious or time-sensitive issues - including possible criminal exposure, child safeguarding, domestic violence, immigration status, imminent procedural deadlines or significant financial loss - the system should reduce the degree of automated guidance and direct the user to a qualified professional, official information service or competent public body.

Meaningful governance also requires continuing evaluation. Accuracy alone is insufficient for a legal education platform. Evaluation should examine source fidelity, use of outdated law, successful referral, user comprehension, the ability to distinguish automated guidance from



professional advice, and whether users can identify the authoritative source after an interaction. Privacy-preserving logs and an accessible mechanism for reporting problematic responses are also important for quality review and correction. These safeguards reflect the accountability and risk-management principles found in the OECD AI Principles and the Council of Europe Framework Convention (OECD, 2024; Framework Convention on Artificial Intelligence, 2024). The precise institutional implementation remains a matter for the project's approved governance arrangements.

5. Application to “Huquq maktabi 2.0”

The “Huquq maktabi 2.0” concept already contains a strong foundation for the proposed framework. It identifies different user groups, envisages lessons, multimedia, tests, quizzes and mini-cases, and contemplates personalization together with AI-supported navigation. Most importantly, the concept treats AI as an auxiliary mechanism for explanation, navigation and recommendation rather than as an independent source of final legal conclusions. That conceptual choice is consistent with the central argument of this article: legal AI in an educational environment can be useful and interactive without being allowed to assume the role of an autonomous professional adviser. The analysis below builds on those approved choices without treating the article's additional safeguards as already adopted project obligations.

A coherent future interface could, for example, distinguish three user-facing functions. A ‘Learn’ mode could focus on explanation, micro-learning, examples and quizzes; a ‘Find the right legal path’ mode could guide users toward verified laws, institutions and official procedures; and an ‘Explain my situation’ mode could permit limited fact-sensitive guidance subject to stronger privacy, source and referral controls. These labels are illustrative rather than proposed amendments to the project's formal terms of reference. Their analytical value lies in showing how the four risk tiers can be translated into visible user choices instead of remaining only an abstract regulatory model.

The project's planned segmentation by age and profession should also be understood as relevant to risk, not only to pedagogy. Younger users may require simpler explanations, clearer warnings against sharing identifying information and earlier referral when a question becomes serious or fact-sensitive, whereas adult or professional users may reasonably receive more detailed regulatory navigation. The proposed framework therefore links personalization to differentiated safeguards as well as to differentiated content. This is a design principle advanced by the article, not a claim that the current application already owes separate legal duties to each user group.

The project concept records technical and functional limitations in the current platform, including weaknesses in analytics, synchronization and user experience. These observations have direct significance for any future expansion into more personalized AI-supported legal guidance. If advanced AI functions are introduced, source management, update procedures, access controls and reliable synchronization should accompany that expansion, because



technical instability can translate into legal unreliability. This article treats those measures as conditions for responsible future functionality rather than as additional obligations retroactively imposed on the existing application.

The concept's reference to anonymized statistics also supports a privacy-preserving approach to evaluation. Aggregate information on course completion, frequently searched topics, unresolved queries and learning outcomes can improve the platform without requiring long-term retention of detailed personal legal narratives. Where analytics are used for research or system improvement, data minimization and separation of identifying information from conversational content should be preferred. The exact technical method will depend on the platform's architecture and approved data-governance arrangements.

The project's publication and research component provides a further opportunity to test the framework empirically. Subject to an approved research design, evaluation should include predefined legal questions, false-premise scenarios, outdated-law examples, multilingual inputs and age-appropriate comprehension tasks. Results should be reported by function and risk level rather than through a single global accuracy figure, because a system may perform well at general explanation while remaining unreliable in personalized or high-stakes settings. Such testing is proposed as a research pathway within the project's scientific agenda, not as an additional grant deliverable created by this article.

6. Implications for Regulation in Uzbekistan

Uzbekistan's 2030 AI Strategy provides a policy basis for developing an AI regulatory framework informed by international experience (Resolution No. PQ-358, 2024). AI-supported legal education is a useful test case because it combines a clear public-interest benefit with identifiable risks of error, over-reliance and misuse of personal data. The analysis does not call for an immediate statute devoted specifically to legal chatbots. It does, however, support several concrete regulatory propositions that could guide sectoral standards, institutional policies and future legislation as national practice develops.

First, public and university legal education platforms should classify the functions performed by AI according to their legal and practical risk before deployment and should revisit that assessment when functionality materially changes. A system used only for general explanation presents a different risk profile from one that later adds personalization, voice interaction, document generation or fact-specific recommendations. Functional change should therefore trigger a fresh review of whether existing safeguards remain adequate.

Second, official legal sources should serve as the primary reference base for automated legal guidance provided through public or university platforms. Uzbekistan already has an established official legal-information infrastructure, which offers a more traceable and legally defensible basis for explanation than unverified web content or opaque model memory. Technical implementation may vary, but the normative principle should remain clear: where the platform



makes a legal proposition, the user should be able to trace that proposition to an authoritative and current source.

Third, high-stakes interactions should trigger human review or authoritative referral within the proposed regulatory model. It is neither necessary nor realistic to define every sentence that constitutes 'legal advice.' A more workable approach is to identify categories of consequence - such as criminal exposure, child safety, imminent procedural deadlines, significant status effects or substantial property interests - in which an autonomous final recommendation creates an unacceptable level of reliance risk. The institutional form of review may differ, but the principle of escalation should be preserved.

Fourth, youth-facing AI systems should be subject to enhanced privacy and interface safeguards. Age-appropriate transparency, data minimization and limits on unnecessary retention of sensitive legal narratives are consistent with international child-centered AI guidance and with the EU AI Act's concern about exploitation of age-related vulnerabilities (UNICEF, 2025; Artificial Intelligence Act, 2024). These sources do not directly govern the Uzbek platform; rather, they provide persuasive regulatory benchmarks for designing domestic safeguards around young users.

Finally, legal literacy should be treated as a component of AI literacy in the legal domain. Users need to know not only that an AI system may be wrong, but also how to verify a legal answer, distinguish an official source from commentary, identify the competent authority and recognize when professional help is necessary. A responsible legal education platform should therefore teach users how to use automated guidance critically. In that sense, the strongest form of legal-AI literacy is not greater dependence on the system, but greater capacity to verify, question and appropriately limit reliance on it.

7. Discussion

A risk-based approach is attractive because it avoids two unhelpful extremes. Treating an educational chatbot as exempt from ordinary legal safeguards merely because it is innovative would overlook the practical consequences of individualized guidance. At the other extreme, treating every personalized interaction as inherently unacceptable would ignore the value of verified explanations and navigation for people who otherwise struggle to access legal information. The relevant question is therefore not whether conversational AI is permissible in the abstract, but what function it is performing and what may happen if that function is wrong.

This functional view also exposes the limits of disclaimers. A notice can help set expectations, but it cannot by itself neutralize reliance where a system solicits detailed facts, produces an individualized conclusion and recommends a specific course of action. The stronger safeguard lies in the interaction design itself - including when the system narrows its answer, cites an authoritative source, signals uncertainty or redirects the user.

Multilingual and jurisdictional performance presents a separate challenge. A platform operating in Uzbekistan may need to handle Uzbek in Latin and Cyrillic scripts, Russian and



potentially other languages. Legal accuracy may not transfer evenly across those settings because training data, retrieval coverage and terminology differ. Evaluation therefore has to be sensitive to language and jurisdiction; a benchmark passed in English cannot simply be assumed to establish the same level of reliability for Uzbek legal terminology.

Personalization also raises a fairness question. Tailoring explanations to age, experience or profession may improve pedagogy, but it should not become a reason to hide legal sources, narrow available remedies or steer some users away from information that remains relevant to them. A useful distinction is between adapting the explanation layer and restricting the underlying legal entitlement: the former may support learning, while the latter requires much stronger justification.

Institutional responsibility remains important even where the underlying model is supplied by a third party. The deploying institution chooses the use case, sources, interface, retention settings and referral logic. Procurement and service arrangements may therefore need to address matters such as safety documentation, data processing, material model changes, incident reporting and testing access. The precise contractual allocation will depend on the technology and the institution; the point is simply that model-provider responsibility does not exhaust the governance question.

8. Conclusion

AI can make legal education more accessible, interactive and responsive, particularly for young people who may find traditional legal materials difficult to navigate. Yet the same conversational qualities that make AI attractive for legal literacy also create a pathway from education to de facto legal advice. The regulatory task is therefore not to decide whether an AI chatbot is simply ‘educational’ or ‘advisory’ in the abstract. It is to govern what the system is doing in each interaction and to match safeguards to the consequences of error.

The article has proposed a four-tier framework: educational explanation, legal navigation, personalized legal guidance, and high-stakes or outcome-determinative legal advice. As the interaction moves across these levels, the framework associates greater legal consequence with progressively stronger source, privacy, transparency, referral and accountability measures. At the highest-risk end, the model favors qualified human review or authoritative referral rather than an autonomous final judgment by an educational platform.

For “Huquq maktabi 2.0,” the framework is broadly consistent with the project’s scientific concept, which treats AI as an auxiliary navigator and explanatory mechanism rather than an independent source of final legal conclusions. The article nevertheless takes a clear normative position: where AI-supported guidance becomes personalized, time-sensitive or capable of materially affecting rights, stronger source, privacy and referral safeguards are justified, and high-stakes interactions should not culminate in an autonomous final legal judgment. These propositions are analytical recommendations for responsible future development; they do not



amend the project's approved tasks or create new contractual commitments for the existing application.

Future research should empirically test the proposed tiers with real users, especially minors and first-time legal-information seekers. Evaluation should measure not only response accuracy but also comprehension, reliance, source verification, referral quality, privacy outcomes and the user's ability to recognize the limits of AI. The long-term objective should be a legal AI ecosystem in which technology expands legal capability while preserving the human judgment, institutional accountability and authoritative sources on which the rule of law depends.

Funding and Acknowledgment

This research was prepared within the framework of the applied research project "Development of the 'Huquq maktabi 2.0' platform and mobile applications aimed at enhancing the legal consciousness and legal culture of youth using artificial intelligence technologies," implemented under State Scientific Project Agreement No. AL-11725099391 (2026-2028). The project is commissioned through the Innovation Development Agency under the Ministry of Higher Education, Science and Innovation of the Republic of Uzbekistan and financed within the state system for funding scientific research and supporting innovation. The article reflects the author's academic analysis and does not represent an official position of the project commissioner, funder, or Tashkent State University of Law.

Declaration of Competing Interests

The author declares no competing interests in relation to this article.

Data Availability Statement

No new empirical dataset was generated for this doctrinal and design-oriented study. The legal and policy materials used are publicly available from the cited official sources. The "Huquq maktabi 2.0" scientific-practical concept is an internal project document used as a case-study source.

REFERENCES

1. American Bar Association Standing Committee on Ethics and Professional Responsibility. (2024, July 29). Formal Opinion 512: Generative artificial intelligence tools. https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf
2. Avery, J. J., Abril, P. S., & del Riego, A. (2024). ChatGPT, Esq.: Recasting unauthorized practice of law in the era of generative AI. *Yale Journal of Law & Technology*, 26, 64. <https://yjolt.org/chatgpt-esq-recasting-unauthorized-practice-law-era-generative-ai>
3. Comoglio, P. (2024). AI Act and online legal advice: Issues regarding access to justice. *Rivista Italiana di Informatica e Diritto*, 6(2), 273–285. <https://doi.org/10.32091/RIID0179>



4. Dahl, M., Magesh, V., Suzgun, M., & Ho, D. E. (2024). Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16(1), 64–93. <https://doi.org/10.1093/jla/laae003>
5. Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, CETS No. 225 (2024). Council of Europe. <https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence>
6. Glesner Fines, B. (2025). AI, UPL, & A2J—Generative AI’s disruptions in the delivery of legal services to low-income individuals. *Washington University Journal of Law & Policy*, 77(1), 126–193. <https://journals.library.wustl.edu/lawpolicy/article/id/9025/>
7. Katz, D. M., Bommarito, M. J., Gao, S., & Arredondo, P. (2024). GPT-4 passes the bar exam. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2270), 20230254. <https://doi.org/10.1098/rsta.2023.0254>
8. Kimbrough, J. L. (2025). Developing lawyering skills in the age of artificial intelligence: A framework for legal education. *Journal of Technology Law & Policy*, 29(1). <https://scholarship.law.ufl.edu/jtlt/vol29/iss1/2/>
9. Law of the Republic of Uzbekistan No. 349-I on Advocacy (December 27, 1996). <https://lex.uz/acts/-54503>
10. Law of the Republic of Uzbekistan No. O’RQ-547 on Personal Data (July 2, 2019). <https://lex.uz/uz/docs/-4396419>
11. Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C. D., & Ho, D. E. (2025). Hallucination-free? Assessing the reliability of leading AI legal research tools. *Journal of Empirical Legal Studies*, 22(2), 216–242. <https://doi.org/10.1111/jels.12413>
12. OECD. (2024). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449; revised 3 May 2024). <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
13. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). (2024). Official Journal of the European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>
14. Resolution of the Cabinet of Ministers of the Republic of Uzbekistan No. 532 on measures to introduce a continuous system of legal education in state educational institutions and to improve the legal culture and literacy of minors (August 27, 2024). <https://lex.uz/uz/docs/-7086385>
15. Resolution of the President of the Republic of Uzbekistan No. PQ-358 on approval of the Strategy for the Development of Artificial Intelligence Technologies until 2030 (October 14, 2024). <https://lex.uz/en/docs/-7158604>



16. Ryan, F., & Hardie, L. (2024). ChatGPT, I have a legal question? The impact of generative AI tools on law clinics and access to justice. *International Journal of Clinical Legal Education*, 31(1), 166–205. <https://doi.org/10.19164/ijcle.v31i1.1401>

17. Tashkent State University of Law. (2026). Scientific-practical concept for the development of the “Huquq maktabi 2.0” platform and mobile applications aimed at enhancing youth legal consciousness and legal culture through artificial intelligence technologies [Unpublished project document].

18. UNESCO. (2022). Recommendation on the ethics of artificial intelligence. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>

19. UNESCO. (2023). Guidance for generative AI in education and research. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>

20. UNICEF. (2025). Guidance on AI and children (Version 3.0): Recommendations for AI policies and systems that uphold child rights. <https://www.unicef.org/innocenti/reports/policy-guidance-ai-children>