


AdvoAI

Legal Intelligence, Simplified.

An AI legal assistant for Uzbekistan. Ask in plain Uzbek, Russian or English — get an answer grounded in the Constitution and the codes, with the article it came from.

advoai.app

Doniyor Xudoykulov · To'rabek Raufov · August 2026



Legal help exists. Access to it doesn't.

A nationwide survey of people who had a legal problem in the previous two years.

8.6%

can afford a lawyer

Fees in the survey ran from under 200,000 to over 2,500,000 UZS for a single matter.

48.6%

go to the mahalla first

Before any state body (28.7%), the other party (23.4%), or a court (10.3%).

1 in 3

blame the law itself

They trace their problem directly to not understanding what the law required of them.

The gap is not a shortage of law. Uzbekistan has published more legislation in the past eight years than in any period before it. The gap is that almost nobody can read it.

The law is public. That is not the same as usable.

Lex.uz publishes every code and statute in full. It offers keyword search — no semantic search, no explanation, no answer to your situation.

The employee

Dismissed without notice. Knows it was wrong. Doesn't know which article says so, or how many days are left to act on it.

The tenant

The landlord is keeping the deposit. The rule exists — four hundred pages into a code they have never opened.

The small business

A supplier contract needs checking. A lawyer's fee is larger than the contract's value.

The newcomer

An investor or migrant who needs the rule in a language they actually read, before signing anything.

Every one of them has the same first question — what does the law say about my situation — and no affordable way to ask it.

Ask in plain language. Get the article.

AdvoAI sits between the citizen and the statute book, and shows its work every time.

01

Ask

A question in ordinary Uzbek, Russian or English. Or upload the contract, notice, or photo of the document you were handed.

02

Retrieve

The question is rewritten into a search query, then matched against the Constitution and the codes of Uzbekistan, held as 10,320 embedded passages.

03


Answer

A plain-language explanation built only from the passages retrieved — naming the article, and linking back to the source on Lex.uz.

What this is not: a general chatbot asked to remember Uzbek law. A general model has no copy of the statute and invents article numbers that read plausibly and do not exist. AdvoAI can only cite what it retrieved, and shows you the passage it used.


Live at advoai.app

Not a mockup. Working product, trilingual interface, installable as an app.



The assistant, in Uzbek

Guest users can ask without an account. Suggested openers cover the questions people actually arrive with — starting a business, inheritance, employment contracts.



A real answer

A user highlights a line from an earlier answer and asks "what is that?" — the follow-up is answered in context, tracing turnover tax through Chapter 66 of the Tax Code and articles 463, 464 and 467.

Grounded, not guessing.

Every answer carries the documents it was built from — and opens them.

Citations are constrained, not generated



The model only sees passages the search returned. It cannot cite an article that was not retrieved, which is the failure mode that makes general chatbots unusable for law.

The source opens in place

Tap a citation and the statutory text appears beside the answer, with a link through to the original document on Lex.uz.

And it says when it might be wrong


Retrieval narrows the room for invention; it does not close it. The product carries that warning on every response rather than hiding it.



Left: the source chips shown beneath every answer. Right: the source panel, opened from a citation — the article text in full, with a link through to the original on Lex.uz.

Two models, two jobs.

A reasoning model decides what to search for. A fast model writes the answer. Most of the quality lives in the step between them.



The corpus

Lex.uz pages are fetched, stripped of interface noise, converted to clean Markdown, split into a parent–child hierarchy, embedded and indexed. 10,320 passages are in production today, covering the Constitution and the major codes. Presidential decrees and the newest amendments are the next ingestion pass.

What already exists.

Two people, six months, no funding, alongside full-time study.

10,320

legal passages embedded and searchable

~15,500

lines of production code across two repositories

30+

API endpoints, ten database tables

3

languages, fully translated interface and answers

12

file types accepted — PDF, DOCX, images, and more

106

commits since the first one on 17 February 2026

Also shipped: account system with email verification and Google sign-in · guest mode with rate limits · installable progressive web app · document and image analysis · an admin console that changes models, limits, prompts and interface text live from the database, without a redeploy.

Six months, from commit to production.

No funding, no team beyond the two of us, and a full university load running alongside it.



The market, from the bottom up.

We would rather show the arithmetic than quote a number nobody can check.

37 million

people in Uzbekistan — the largest population in Central Asia, with roughly 28 million online.

91.4%

cannot afford a lawyer today. That is not a market we take from firms; it is demand no one currently serves.

~\$5 / month

is the intended Pro price — under one percent of median monthly income, and a fraction of one consultation.

100,000

paying users would be roughly \$6M in annual revenue. Reaching that is the hard part; the arithmetic is not.

Beyond Uzbekistan: Kazakhstan, Kyrgyzstan and Tajikistan share a post-Soviet civil-law structure. The corpus changes; the pipeline does not.

Affordability figure from the UNDP assessment cited earlier. Population and connectivity are public estimates. The revenue line is illustrative arithmetic at a stated price, not a forecast.

Free for citizens. Revenue from the professional side.

The free layer is not charity and not a loss leader — it is where the demand signal is generated.

Free

Now

Guests get a few questions without an account; registered users a daily allowance. Every limit is tunable from the admin console.

Pro

~\$5 / month

Unlimited questions, deeper document analysis, priority response. Priced against a single consultation, not against software.

Marketplace

Per qualified lead

When a case needs representation, it goes to a licensed advocate — already classified by area of law. The lawyer pays for a client with clear intent, not a cold contact.

API

Per query

Banks, insurers, HR platforms and government portals embedding grounded legal answers in their own products.

Citizens generate the questions. The questions reveal which cases need a lawyer. That flow is what the professional side is paying for — which is why the consumer tier stays free.

Where we draw the line.

An AI answering legal questions in a jurisdiction with a licensed advocacy profession has to be explicit about what it is not.

Information, not representation

AdvoAI explains what the law says. It does not act for anyone, appear before a court, or replace an advocate. Anything requiring representation is routed to a licensed professional.

Sources shown, always

Every answer names its article and links the source document. The user is always one tap from checking us.

Fallibility stated in the product

Each response carries a visible notice that the assistant can be wrong and important matters should be verified — not buried in terms of service.

Privacy and access control

Role-based limits, an opt-out that redacts a user's data from review, and account controls built in from the start rather than retrofitted.

The open question we want help with: exactly where the boundary sits between legal information and legal advice under Uzbek rules on advocacy. We have taken the conservative reading. We would like to be sure of it before we scale.

Who is building it — and what we need.

Doniyor Xudoykulov

AI & Data

Data Science at School 21. Built the retrieval pipeline, ingestion and backend. Prior work in RAG systems and applied machine learning.

To'rabek Raufov

Product & Engineering

Co-founder and technical lead at iTravel. Full-stack engineer; built and shipped tourfixer.uz. Owns the frontend and product surface.

The ask

Regulatory clarity

An introduction to counsel who can tell us where the advice line sits, before we scale rather than after.

Distribution

Warm introductions to law firms, legal-aid bodies and platforms that would embed grounded legal answers.

Adversarial testing

Domain experts willing to try hard to break the assistant, so we find the failures before users do.

Capital: \$40,000 covers twelve months — inference and infrastructure, expanding the corpus to decrees and current amendments, paid legal review of answer quality, entity registration, and modest founder stipends.