

A Personal AI Workstation for Every Lawyer

A concept paper on a national Legal Copilot ecosystem for Ukraine

1. The Problem

Ukraine's legal system is operating under a load that no system would withstand even in peaceful times.

According to the High Council of Justice, as of late 2025, roughly 4,350 judges are actually serving against an authorized complement of 6,600 — more than 2,250 positions are vacant ([sud.ua, 2025](#)). Court caseloads are growing across all jurisdictions — by at least 8% per year — and judges themselves estimate the workload at 5–10 times the European average ([sud.ua, 2026](#); [zmina.info](#)). State legal services, local self-government bodies, ministries and law-enforcement agencies are in the same position: fewer people than tasks, and a rising cost of error.

At the same time, Ukraine has entered the most intensive law-making period in its history. The screening of national legislation for compliance with EU law was completed in September 2025; on 15 June 2026 the first negotiation cluster, "Fundamentals," was officially opened ([Council of the EU](#)). Ahead lie the harmonization of tens of thousands of pages of the *acquis communautaire*,

justice reform under the €50 billion Ukraine Facility ([Council of the EU](#)), and the post-war rebuilding of institutions.

The bottom line is simple: the volume of qualified legal work the country must perform in the coming years exceeds the physical capacity of its existing specialists many times over. Hiring the missing lawyers is impossible — they do not exist. One path remains: radically raising the productivity and quality of everyone already working.

2. Why the Old Approach No Longer Works

The state's standard answer to such a problem is well known: announce a major digital reform, gather requirements, run a tender, commission a single centralized system.

This path predictably leads to the same result:

- ♦ requirements gathering averages the needs of thousands of different specialists into a lowest common denominator;
- ♦ approvals and procurement procedures take years;
- ♦ building a heavyweight system costs tens of millions;
- ♦ by the time it is deployed, the system is already two or three technology generations behind;
- ♦ every change goes through the contractor and a new budget cycle;
- ♦ the specialist receives a tool designed "for everyone" — that is, for no one in particular.

This is not a Ukrainian peculiarity but a global pattern. The OECD, having analyzed 200 cases of AI use in government functions, concludes that most initiatives get stuck at the pilot stage and never reach effect at scale ([OECD](#),

[Governing with Artificial Intelligence, 2025](#)). The problem is not the models — it is the organizational model of adoption.

In a field where the technology refreshes every six months, designing "a single perfect product for the decade" means designing guaranteed obsolescence.

3. What Has Changed in Technology

Over the past three years, three shifts have occurred that together open a window of opportunity.

First shift: language models learned to work with legal text at a professional level. As early as 2023, GPT-4 passed the US bar exam at the ~90th percentile ([Katz et al., 2024](#)). By 2025–2026, independent blind evaluations showed specialized legal AI tools performing research tasks at 74–78% versus a 69% lawyer baseline ([Vals Legal AI Report, October 2025](#)). Current flagship models — GPT-5.5 ([OpenAI, April 2026](#)), Claude Opus 4.7 ([Anthropic, April 2026](#)), Gemini 3 ([Google, 2025](#)) — reach the level of domain experts on well-specified professional tasks ([GDPval, OpenAI](#)).

Second shift: agentic working environments appeared. OpenAI Codex ([documentation](#)), Claude Code ([documentation](#)), and the open-source Gemini CLI ([GitHub](#)) and opencode ([opencode.ai](#)) are no longer chatbots. They are environments in which the model reads project folders, executes multi-step processes, uses tools, keeps memory, and follows project instructions and specialized roles. The AGENTS.md project-instruction format has become a cross-vendor standard: it is used by more than 60,000 open projects, and its specification has been handed to the Linux Foundation ([agents.md](#)). The Model Context Protocol has standardized how models connect to external data and tools ([MCP specification](#)). This is ready, industrially proven infrastructure: it does not need to be developed — it needs to be configured for legal work.

Third shift: data and tools can now be connected safely at the level of an individual workstation. You do not need to build a megasystem to give a lawyer a model connected to current legislation, case materials and verification scripts. A properly organized local configuration is enough.

Ukraine has approached this window better than most countries. Diia has surpassed 20 million users, and the country ranks 5th in the world in the UN Online Service Index ([UN E-Government Survey 2024](#)). In September 2025, Diia.AI was launched — the world's first national AI agent that does not merely advise but actually delivers government services ([Ministry of Digital Transformation](#); [European Commission](#)). The Unified State Register of Court Decisions holds more than 120 million decisions in open access (reyestr.court.gov.ua), and the Verkhovna Rada provides a machine-readable legislation API (data.rada.gov.ua).

The technological preconditions are in place. The barrier is not the models.

4. How to Think About a Language Model

For a lawyer, a judge or the head of an institution, it is useful to picture a language model as an **intelligent engine for working with text and reasoning** — a reasoning engine. This engine can:

analyze texts, identify logical connections, map facts to legal rules, formulate legal arguments, find contradictions, prepare draft documents, consider alternative legal positions, model the consequences of decisions, and summarize large bodies of information.

This capability emerged from training on vast corpora of text — including legislation, court decisions, academic publications and professional literature. That is why modern models possess a certain "legal intuition" by default: they

understand the structure of legal argument, the language of normative acts, and typical legal constructions.

But precision of wording is essential here. The model does not "know the law" and does not replace the lawyer.

A language model possesses a powerful general-purpose mechanism for language and reasoning, but it needs current, verified and properly organized data.

An engine without fuel, instruments and a driver is not transport. A model without data, instructions and a qualified user is not a legal tool.

5. Why Plain ChatGPT Is Not Yet a Legal Copilot

Access to a strong model does not by itself create a legal assistant. The reasons are fundamental:

1. **Outdated knowledge.** The model was trained on a snapshot of data that was already stale at release. During EU integration, Ukrainian legislation changes weekly.
2. **Mixed versions and sources.** The training data contains different versions of the same acts — the model can imperceptibly "splice" an old and a new version together.
3. **No reliable temporal anchoring.** The model may fail to establish which rule was in force on the date of the events — and for a lawyer this often decides the outcome of the case.
4. **No materials of the specific case.** The model does not know the facts and the evidence until it is given controlled access to them.
5. **The probabilistic nature of its answers.** The model produces the most plausible text; it does not retrieve guaranteed truth from an official register.

6. **No built-in accountability.** The model does not itself determine when mandatory human review is required.

These are not theoretical risks. Public LLMs hallucinate on legal questions in 58–82% of cases ([Dahl et al., Journal of Legal Analysis, 2024](#)). Even commercial legal RAG tools marketed as "hallucination-free" return unreliable answers on 17–33% of queries ([Stanford RegLab, Hallucination-Free?, 2024](#)). The database of court cases involving fabricated AI content stands at **1,733 cases across 40 jurisdictions** as of 9 July 2026 — against ~370 in October 2025 ([AI Hallucination Cases Database](#)). In the first quarter of 2026 alone, US courts imposed more than \$145,000 in sanctions for fabricated citations ([NPR, April 2026](#)).

The Supreme Court of Ukraine has already produced its first case law on the matter: in a ruling of 8 July 2025, the Commercial Cassation Court within the Supreme Court held that references to AI answers cannot be used to refute the conclusions of a court ([sud.ua](#)).

The conclusion: a complete Legal Copilot is not a model. It is a model **plus** data, instructions, tools, verification mechanisms and rules of interaction with the user.

6. The Formula for a Complete Copilot

***A strong language model*

- ◆ *current legal data*
- ◆ *the materials of the specific matter*
- ◆ *professional instructions*
- ◆ *verification tools*
- ◆ *a qualified user = a personal legal copilot.***

Every term is mandatory. Remove current data — and you get confident answers based on repealed rules. Remove the matter materials — and you get abstract reasoning. Remove the instructions — and you get social-media style instead of procedural discipline. Remove verification — and you get case number 1,733 in the hallucination database. Remove the qualified user — and you get all of the above at once.

The main effect arises from the synergy of a general-purpose intelligent engine with current, structured, properly delivered data.

7. The Architecture of Data Layers

The data a legal copilot works with naturally falls into five layers with different access regimes.

The public normative layer. The Constitution, codes, statutes, secondary legislation, international treaties, case law, official interpretations, classifiers, open registers. The state's task is to publish this layer not only as web pages and PDFs but as machine-readable data: a stable identifier, adoption date, period of validity, version history, structural markup, citation metadata. International standards for this exist: Akoma Ntoso / LegalDocML ([OASIS](#)) and the European Legislation Identifier ([ELI](#)). The United Kingdom publishes its entire statute book through an open API ([legislation.gov.uk](#)), the United States publishes its Code in USLM XML ([uscode.house.gov](#)), and Germany has digitized its entire legislative cycle in LegalDocML.de ([E-Gesetzgebung](#)). For an EU candidate country, moving to ELI-compatible markup is not optional — it is part of the integration agenda.

The institutional layer. Internal regulations, methodologies, templates, policies, the institution's accumulated practice, standard legal positions. This is the organization's institutional memory — today it lives in the heads of departing employees and in unnavigable network folders.

The confidential layer. The materials of a specific case, personal data, evidence, attorney-client privilege, restricted official information. Here the architecture must guarantee access control, logging of every request, and the option of local or isolated processing. The reference point is the reform of Ukraine's personal-data legislation toward GDPR standards (draft law [No. 8153](#)).

The personal layer. The specialist's individual working memory: templates, analytical approaches, personal instructions, notes, a history of improvements.

The team artifact layer. Structured work products exchanged by members of a working group. This layer is central to the entire concept — and it deserves a separate discussion.

8. The Personal Agentic Workstation

The atomic unit of the proposed infrastructure is not an application or a portal but the lawyer's **personal agentic workstation**.

Technically, it is a configured agentic environment (OpenAI Codex, Claude Code, open-source equivalents, or a locally hosted government build) that can:

- ◆ read the documents and folders of a specific case;
- ◆ query permitted legal sources — registers, legislation APIs, case law;
- ◆ run search and comparison;
- ◆ prepare draft legal documents;
- ◆ verify references and quotations against primary sources;
- ◆ maintain the specialist's working memory;
- ◆ use specialized roles and skills;
- ◆ keep a history of changes;

- ♦ assemble team artifact packages and accept packages from others;
- ♦ work with local and remote data sources.

The concept's central practical claim: the workstation does not need to be developed. Its software foundation already exists as ready, industrially proven infrastructure — the agentic environments OpenAI Codex, Anthropic Claude Code and their open-source equivalents. "Creating" a lawyer's workstation is not a software project but a set of text files on top of this infrastructure: instructions, roles, skills, templates, verification rules. That is why deployment can start today — with no tender, no contractor, no custom product development.

The key property is **configuration without programming**. Modern agentic environments are configured with text files: project instructions (AGENTS.md), roles, skills, process templates, citation rules, verification procedures. This means:

Professional lawyers should not program the model — they should formalize their own practice in clear instructions, examples, templates and control rules.

Legal expertise becomes a reproducible digital asset. An instruction on "how we review a municipal-property lease agreement," written by the department's best lawyer, becomes a file that works in every employee's workstation — and survives staff rotations.

The second key property is **provider independence**. The workstation is built on open formats (Markdown, JSON, YAML) and standard protocols (MCP, AGENTS.md). Switching the model or the vendor replaces the engine, not the whole system. For the state, this is the strategic answer to the risk of dependence on foreign suppliers.

9. From Documents to Artifacts

Here begins the most important part of the concept — the move from individual efficiency to team efficiency.

In traditional teamwork, the object of exchange is the final document: a memo, an opinion, a memorandum, a draft decision. The document shows the result but hides the process. The recipient cannot quickly establish: which sources were used; which queries were run; which documents were discarded; which versions of the legislation were analyzed; which alternative hypotheses were considered; which checks were performed; why this particular conclusion was reached.

Review comes down to reading the text and redoing a substantial part of the work. A finished document is a conclusion with the proof erased.

With the arrival of AI, this asymmetry became dangerous: **text production accelerated by an order of magnitude, while verifiability remained manual**. A team in which everyone sped up individually but still exchanges work the old way has not become faster — it has become riskier.

The proposed shift: the final document ceases to be the sole object of exchange and becomes one of the **artifacts** in a work package.

The document is the interface. The work package is the legal object.

The work package (in VeriLex terminology, the Legal Work Package, LWP) includes: the final text; a source registry with stable links; extracted fragments of primary sources; a fact matrix with statuses (asserted / established / disputed / inferred); a matrix of applicable rules; a tree of legal issues; identified contradictions and alternative hypotheses; a log of key decisions; verification protocols; human comments; a confidence assessment; a list of open questions; a change history; machine-readable metadata.

All of this comes in open, portable formats: Markdown, JSON/JSONL, YAML, CSV, Mermaid graphs, Git versioning. No proprietary containers.

The team receives not only the answer but the **reproducible context** from which that answer was formed.

10. The Provenance Chain and Process Verification

Every substantive conclusion in a work package carries a verifiable provenance chain:

Source → fact extraction → interpretation → mapping → intermediate conclusion → final claim.

An important caveat: this is not about disclosing the model's hidden internal reasoning — that is legally useless. It is about recording **verifiable working traces**: sources, facts, intermediate propositions, mapping rules and completed checks. This is the same logic as the W3C PROV data-provenance standards ([PROV-O](#)) and the culture of reproducible research ([ACM Artifact Review](#)), transferred to legal work.

The recipient's agentic system can check the package automatically: completeness of sources; currency of legislative versions; correctness of dates; fidelity of quotations to primary sources; logical coherence of the argumentation; missed counterarguments; changes between versions; compliance with the internal research standard.

The team moves from asking "is this text written correctly?" to asking "how was the information it rests on collected, verified and analyzed?"

Two principles keep the system from degenerating into bureaucracy:

No source — no strong conclusion. The agent proposes — the human decides. Always.

And the principle of sufficient transparency: three levels of artifact detail. **Basic** — for everyday work (text, sources, key facts, a brief verification log, open questions). **Extended** — for complex research (plus the argument structure, legislative versions, alternative positions, a relationship graph). **Audit-grade** — for decisions with a high cost of error (plus the full change history, model identifiers, checksums, participant roles, human confirmation). The goal is not maximum logging but sufficient reproducibility of the result.

11. Team Memory and Reuse

Structured artifacts solve another chronic problem — the loss of institutional memory.

When a similar task arises again, the agentic system can: find the earlier research; determine which parts of it are still current; check how the legislation has changed since the previous analysis; reuse confirmed sources; identify discrepancies; build an updated version and indicate exactly what has changed.

Every completed project becomes not an archived document but a reusable intellectual asset. For Ukrainian public authorities — where staff turnover is high and the load on those who remain keeps growing — this is not an "optimization"; it is a survival condition for quality.

The manager, in turn, gains a new class of management information: a short conclusion plus a map of its grounds, a list of risks, alternative scenarios, a degree of confidence, a list of unconfirmed assumptions, and the questions that require a human decision. Facts, interpretations, assumptions and recommendations stop being fused into an inseparable text.

12. An Open Repository of Base Workstations

For the system to take off, the entry threshold must be minimal. The instrument is a public repository of Legal Copilot starter configurations: base instruction files, standard roles, ready-made skills, document templates, verification checklists, citation rules, source-connection guides, a team artifact standard, training projects, test cases, security recommendations.

The starter set is 5–10 configurations: an assistant for a lawyer in a public authority; a judge's assistant; an advocate's assistant; a legislative drafter's assistant; a contracts assistant; an administrative-law assistant; a case-law assistant; an assistant for the legal review of normative acts; a municipal lawyer's assistant; an educational workstation for a law student.

The user scenario: download the right configuration, install it locally, connect your agentic environment — and get your first working workstation in an hour, not in a year.

Note: the repository contains not a single line of custom-built code. It is a complete "distribution" of the national system, consisting of instructions, roles, skills, standards and training courses on top of already existing agentic environments. That is precisely why what stands between the launch decision and the first working workstation is the publication of files, not a development cycle.

Then the main thing begins: **self-development**. As their skills grow, users improve the instructions, templates, skills and control procedures to fit their own practice. The workstation becomes a reflection of the professional experience of the individual specialist. The best solutions flow back into the repository — and thousands of users independently develop the shared ecosystem. This produces a distributed-development effect unattainable by any centralized contractor.

13. Education and Mentorship

Technology is adopted at the speed at which people learn, not at the speed at which software installs.

Universities and professional schools — law faculties, the National School of Judges, civil-service training centers, bar schools — take on an active role in this model: developing legal skills for agents; testing model quality on Ukrainian legal tasks; preparing teaching cases; researching errors and limitations; training students; shaping the standards of legal AI literacy. Knowing how to organize the work of a legal agent and hand colleagues a verifiable artifact should become a baseline professional competence — as working with registers is today.

Mentors are the second circuit. A mentor is neither a programmer nor a system administrator. The mentor's task is not to configure the system for the user but to teach the user to develop their own workstation: to formulate tasks, describe processes, create instructions, define sources, introduce result verification, set up the production and acceptance of team artifacts — and gradually hand over full control. Commercial models for mentorship (hourly work, certified programs, institutional support, expert-support subscriptions) make this role a sustainable profession rather than volunteering.

The first official steps have already been taken: in July 2025 the Ministry of Digital Transformation, together with the Ministry of Justice, the State Judicial Administration and the Ukrainian Bar Association, issued recommendations on the safe use of AI by legal professionals (thedigital.gov.ua).

14. A National Model for Ukraine

Why is Ukraine the country that can implement this model first?

A proven capacity for digital breakthroughs. A country that launched Diia.AI — the world's first national AI agent for government services ([Cabinet of Ministers of Ukraine](#)) — knows how to turn concepts into working services faster than most EU states.

The institutional framework already exists. The WINWIN strategy to 2030 ([Ministry of Digital Transformation](#)), the White Paper on AI regulation oriented toward the EU AI Act ([Ministry of Digital Transformation, 2024](#)), and the draft AI Development Strategy to 2030 with the goal of entering the top 3 countries for AI integration in the public sector ([Digital State](#)).

Open legal data. More than 120 million court decisions in the open register, machine-readable datasets with daily updates ([data.gov.ua](#)), the parliament's REST API ([data.rada.gov.ua](#)), and 40,000+ datasets on the national portal. Few EU countries can show a comparable open legal corpus.

A mature LegalTech community. Opendatabot, YouControl, LIGA ZAKON, ZakonOnline, Vkursi, Clarity Project — a market that has spent a decade building products on top of open government data ([Mezha.Media](#)).

The EU-accession window. Harmonization with the *acquis* is an enormous volume of legal work — but it is also a source of funding (Ukraine Facility: more than €29.5 billion already disbursed against reform delivery) and a ready-made system of standards: ELI for identifying legislation, the EU AI Act's requirements for high-risk systems in justice (taking effect on 2 December 2027 after the Digital Omnibus postponement — [Covington, 2026](#)), the CEPEJ principles ([Council of Europe](#)) and the UNESCO guidelines on AI in courts ([UNESCO, December 2025](#)).

The state's role in the proposed model is deliberately limited. The state does not design "a single perfect Legal AI product." It centralizes only what no one else can: machine-readable law; data and artifact standards; security and certification; base configurations; training; pilots; interoperability. Everything

else — a multitude of quality solutions — will grow faster in a distributed ecosystem of practitioners, universities and the market.

15. Pilot Projects

The launch is not a multi-year reform but a controlled chain reaction. None of the eleven pilots requires custom software development: the agentic environments are ready and available (commercially or as open source), the data is open, and the configurations are files. The technical start is possible immediately after agreements with the institutions are signed. Eleven candidates for national pilots:

1. **A workstation for a local self-government lawyer** — typical community requests, contracts, municipal property.
2. **An assistant for analyzing draft normative acts** — review of draft laws, detection of conflicts, mapping against the *acquis*.
3. **A judge's assistant for structuring case files** — chronologies, evidence maps, fact matrices.
4. **A civil servant's assistant for preparing legal responses** — citizens' appeals, information requests, clarifications.
5. **A university workstation for students and faculty** — teaching cases built on real open data.
6. **An assistant for harmonizing legislation with EU law** — mapping Ukrainian rules to directives, tracking progress.
7. **A workstation for case-law analysis** — the 120 million decisions of the Unified State Register of Court Decisions as a corpus for verifiable analysis.
8. **An assistant for state legal services** — claims and litigation work, representation.

9. **A tool for reviewing contracts and regulations** — typical risks, compliance with legislation.
10. **A training and certification system for Legal AI mentors** — the talent pipeline of the entire ecosystem.
11. **A pilot of a working group that has fully switched to exchanging structured artifacts** — the key pilot, proving the team effect.

Detailed pilot passports (target user, problem, data, configuration, artifact format, risks, success criteria, timelines) are in the [pilot launch plan](#).

An indicative first-year trajectory (a working hypothesis that requires refinement): **months 0–3** — the open repository, 5–10 configurations, a minimal artifact standard, training materials, selection of pilot institutions. **Months 3–6** — workstation rollouts, a team artifact-exchange pilot, mentor training, feedback collection. **Months 6–12** — scaling, new specializations, interoperability standards, a national network of practice.

16. Risks and Limitations

The concept does not stand without an honest conversation about risks. Sixteen principal ones — with mitigation mechanisms.

Hallucinations and errors. Mitigation: mandatory binding of conclusions to sources (no source — no strong conclusion), automated checking of quotations against primary sources, human confirmation of legally significant decisions.

Outdated sources. Mitigation: machine-readable data with version validity dates, automatic currency checks on every use.

Leaks of confidential information. Mitigation: a layered access architecture, local processing of the confidential layer, redaction before calls to external models, logging.

Dependence on foreign suppliers. Mitigation: a provider-agnostic architecture, open formats, the ability to switch models, development of local and open models.

Model opacity. Mitigation: transparency is provided at the level of the process (artifacts, the provenance chain), not the model's internals.

Automated bias. Mitigation: testing on Ukrainian legal tasks, critique packages with mandatory counterarguments, model diversity.

Insufficient user qualification. Mitigation: the mentorship institution, educational programs, certification, base configurations with protective rules.

Absence of common standards and fragmentation. Mitigation: an open artifact standard from day one, interoperability validators, a coordinating center.

Overtrust in answers. Mitigation: mandatory confidence assessments and open-question lists in every package, a culture of "the agent proposes — the human decides."

Accountability problems. Mitigation: legally significant actions are technically available only to a human; every decision has a specialist who confirmed it.

Institutional resistance. Mitigation: voluntary pilots with measurable effect, not forced top-down deployment.

Excessive logging and a false sense of control. Mitigation: three levels of artifact detail; the usefulness metric for logs is their participation in real checks and decisions.

Confidential data ending up in artifacts. Mitigation: automated package checks before transfer, anonymization rules, layer separation.

The full risk register, together with skeptics' counterarguments, is in a [separate publication](#).

17. A Plan for the First Steps

Igniting the chain reaction requires neither billions nor the development of a software product: the technological foundation has already been built by the global industry and is used by millions. What is needed:

1. several high-quality base workstations (configurations, not development from scratch);
2. an open repository with clear documentation;
3. access to machine-readable legal data (a significant part is already open);
4. a minimal team artifact standard;
5. short practical courses and demonstration cases;
6. a first-wave mentor network;
7. security standards for working with the confidential layers;
8. a venue for exchanging skills;
9. a pilot deployment program in 3–5 institutions;
10. effect measurement: speed, quality, verifiability — before and after.

Each item is achievable by a small team with the support of one or two institutional partners and international technical-assistance programs (EU Project Pravo-Justice, EU4DigitalUA and similar).

18. Conclusion

Seven answers this concept gives:

Why now? Models have reached professional level on legal tasks, agentic environments have become the standard, and Ukrainian legal data is already open — the window of opportunity opened in 2024–2026. The workstation infrastructure will not "appear in the future" — it is ready to use now.

Why is one model not enough? The model is an engine, not transport. Without current data, instructions and verification it produces plausible text, not verified text — 1,733 court cases about hallucinations are the proof.

How do you connect models, data and expertise? Through the personal agentic workstation: the formula "model + data + matter materials + instructions + verification + user."

Why not a single state SaaS? Centralized development averages, ages and creates dependence. The state should build not a product but the conditions: data, standards, security, training.

How do you move from personal efficiency to team efficiency? Change the unit of exchange: not the final document but a structured work package with a verifiable provenance chain for every conclusion.

Why are artifacts better than documents? Because a document shows only the result, while a package shows the result together with its grounds, checks and open questions. Review turns from redoing the work into auditing the process.

What can Ukraine do in a year? A repository, an artifact standard, 5–10 configurations, a mentor network, 3–5 pilots with measured effect — and the world's first national model of distributed legal AI infrastructure.

There is no need to start with a national megasystem. Start with ten quality workstations, open data, a shared artifact standard, and a community of professionals.

And the start can be today. The entire software infrastructure already exists — what remains is to publish in an open repository the things that turn it into legal workstations: instructions, skills, roles and training courses.

Then the thing that has always worked better in Ukraine than directives will take over — the distributed energy of people who have been given a good tool.

*Author of the concept and the text: **Sergej Avdejcik** — creator of the VeriLex project (verifiable AI workflows for legal teams), researcher of agentic workflows in legal practice.*

Evidence base: [confirmed claims and the full source table](#). The argumentation map is part of the project's working materials.