

**COLORADO SUPREME COURT  
LEGAL TECHNOLOGY ADVISORY COMMITTEE  
AGENDA**

June 25, 2026, noon MST  
Court of Appeals Full Court Conference Room and via Webex

Webex link:

<https://judicial.webex.com/judicial/j.php?MTID=m8ec891dd66d4f8d2e572624d23923c07>

---

1. Call to order.
2. Approval of minutes of the May 28 meeting.
3. Discussion of the first draft of the consolidated document and the most efficient process for providing comments.
4. Discussion of the sections that still need to be drafted.
  - a) Universal principles and key recommendations.
  - b) Index.
  - c) Glossary.
  - d) FAQs.
  - e) What else?
5. Thoughts on the report of the Colorado Bar Association's AI Task Force.
6. Next steps.
7. New business.
8. Adjournment.

Upcoming meeting date: September 10, 2026 (Supreme Court Conference Room).

Judge Lino Lipinsky, Chair  
Colorado Court of Appeals  
[lino.lipinsky@judicial.state.co.us](mailto:lino.lipinsky@judicial.state.co.us)

## Phase A — Per-Paper Structured Analyses

### Colorado LTAC AI Guidance Document — Synthesis (Phase A only)

Documents read. Controlling authority: 01 (Chief Justice Márquez's charge / project plan) and 01a (Feb. 10, 2026 minutes establishing the seven subcommittees). Substantive drafts analyzed: 03 Fundamentals, Transformational Aspects of AI Technology, 04 Generative AI Resources, 05 Professional Ethics & UPL, 06 Self-Represented Litigants, 07 Generative AI and Legal Professionals.

Not located / flagged for confirmation. The charge mapping anticipates 06 in two parts (pt 1 & pt 2). Retrieval returned the SRL material abundantly as one body (sections I–IV: Guiding Principles, Standards for Responsible Development, Risks, Verification, Steering, Infrastructure, Two-Tier Justice, Responsibility of the Courts, Privilege Question), but no separately-titled "Part 2" surfaced. This analysis treats the retrieved SRL material as a single paper. *Please confirm against your folder whether a distinct Part 2 exists*; if it does, it needs its own A1–A7 pass. No content has been invented to fill the gap.

Citation anchor convention. [doc short-title, § "heading", "short phrase"]. Page/line resolution is a downstream human step against the Word files. All quoted phrases are short locator strings, not block quotes.

---

### 1. Fundamentals of AI Technology (03)

*Analyzed from the standpoint of the Fundamentals subcommittee draft (Moriarty/Kaminski/Banks).*

A1. Core conclusions / guidance. Non-technical users should first "play with it," but hands-on use is insufficient because not every failure is observable from experimentation; technical understanding of the architecture is necessary because "the power and dangers of generative AI flow directly from its architecture." The draft's operative stance is *cautious, skeptical optimism* and a human-in-the-loop posture. It explains LLMs as transformer neural networks grown by gradient descent ("a giant mathematical equation"), conceptualized via the "shape/blanket over poles" metaphor, then catalogs error states and limitations: hallucination, deception/false reasoning in latent space, context-window limits, tokenization quirks, jailbreaking/representation steering, and the non-human, unaccountable nature of the model.

A2. Key supporting premises / evidence. Architecture account (neurons, attention/transformer blocks, RLHF "warping the shape"); empirical anchors via footnotes — legal-hallucination rates of 58–88% in early models (Dahl/Magesh/Suzgun/Ho), "lost in

the middle" context degradation (Liu), under-trained-token effects (SolidGoldMagikarp), representation steering raising a "blackmail" pattern ~22%→~72%, and interpretability limits ("~3% of features" understood). The reasoning is that capability and risk are two faces of the same architecture, so verification is mandatory.

A3. Accuracy & quality assessment. Technically strong and unusually well-sourced for a guidance document; claims are tied to named papers. Watch items: (a) heavy reliance on perishable specifics — "around one million tokens," "thousand neurons ... hundred layers deep" — true today, dated quickly. (b) The "deception/false reasoning" framing (latent-space answer then post-hoc rationalization) is defensible but strongly stated; ensure Phase B keeps it proportionate. (c) Citation-hygiene blemish: footnote to Paulo & Belrose contains the parenthetical "as discussed in the corpus's feature-consistency synthesis," which reads like a drafting/research note leaked into a citation [03 Fundamentals, footnotes, "feature-consistency synthesis"] — flag for cleanup. No fabricated authorities detected.

A4. Priorities — what this paper most wants preserved (inference). *Inferred from repetition and placement*: (i) that LLMs are a *non-human* intelligence not to be anthropomorphized ("falls into a silent coma"); (ii) that errors can be invisible, so a human-in-the-loop is non-negotiable; (iii) the architecture→risk causal chain as the conceptual "why" the rest of the document depends on. Textual basis: the framing sentence "The power and dangers ... flow directly from its architecture" and the recurring return to error states.

A5. Target audience. *Primary*: legal professionals and judicial officers as non-technical readers needing conceptual grounding (inferred from "non-technical users," lawyer/assistant analogies). *Secondary*: any reader of later sections — it is positioned as foundational. *Drift flag*: occasional jumps to research-grade detail (superposition, sparse autoencoders) momentarily address a technical audience rather than the lay practitioner; signal or relegate to footnotes.

A6. Charge coverage. Subcommittee charge (definition, technical foundations, how output is produced, capabilities & limitations): fully answered. Five objectives: integrity & competent service supported indirectly (grounds verification duty); innovation supported (demystifies tech); confidentiality and access-to-justice only touched. Maps cleanly to the charge's "Fundamentals of AI technology" core issue.

A7. Citation anchors (for Phase B). [03 Fundamentals, § "The Architecture of a Large Language Model", "flow directly from its architecture"]; [03, § "Hallucinations", "confidently present false information"]; [03, § "Limitation of the Context Window", "lost in the middle"]; [03, § "Model Manipulation and Jailbreaking", "turning it up or down"]; [03, § "The Nature of LLMs and Their Error States", "silent coma"].

---

## 2. Transformational Aspects of AI Technology

*Analyzed from the standpoint of the Transformational Aspects subcommittee draft (Guillon/Dorancy/Moriarty/Riehl).*

A1. Core conclusions / guidance. Generative AI's legal use cases are "innumerable"; the draft maps high-level capabilities (drafting/editing, document-set analysis, research, fact development, e-discovery, knowledge management, argument/decision analysis, case management, communication, translation, access to justice, education), then dramatizes them through five user stories (Litigator, In-House/Transactional Counsel, Regulatory Lawyer, Judge, Self-Represented Litigant) before forecasting "Future Uses" (agentic workflows; human in/on/in-command/out of the loop; increased scale; AI assisting decision-makers; A2J). The throughline: AI augments but does not replace human judgment.

A2. Key supporting premises / evidence. Concrete worked examples (e.g., a litigation review "first two weeks" compressed to "ninety minutes"; an uncontested divorce filed for "less than \$300"); the "treat any model output as a 75% answer" heuristic ("the user owns the final 25%"); a tiered-reliability model for pro se users (general chatbot → vetted legal-aid content → court self-help certification).

A3. Accuracy & quality assessment. The user-story vehicle is vivid and concrete (valued per the soft organizing instruction). Principal flag — charge-scope drift: the subcommittee was expressly charged with *capabilities, not misuse* [01a, § "Transformational Aspects", "rather than on how it has been used and misused"], yet the draft embeds an extensive risk apparatus — inline "Risk" callouts (hallucinations, fake/deepfake evidence, single-shot errors, surrendering judgment, garbage-in/garbage-out, poor legal judgment) and a full "Institutional Problems" section (frontier-lab capture, erosion of the apprenticeship pipeline). This is good material but sits across the 03↔Transformational and risk-home tensions the charge framework flagged; Phase B must decide its home. Secondary flags: the "75%" figure is a useful illustration, not a measured rate; the claim that LLMs "can be more accurate than human judges" where law is settled appears without a retrieved citation [Transformational, § "Assisting Decision Makers", "more accurate than human judges"] — needs support or softening. Dated specifics (model behaviors, dollar figures, tool names) are perishable illustrations.

A4. Priorities — what this paper most wants preserved (inference). *Inferred from form and emphasis*: (i) the user-story device itself — concreteness over abstraction; (ii) breadth of use cases across litigation, transactional, regulatory, and judicial work; (iii) the augment-

not-replace principle and tiered-reliability ladder for pro se users. Textual basis: the deliberate "zoom in on specific users" framing and the repeated Benefit/Risk pairing.

A5. Target audience. *Primary*: legal professionals and judicial officers who "don't yet see why the tech is useful or dangerous" (the imagination-firing role assigned in the reading order). *Secondary*: self-represented litigants (the SRL story) and the public. *Drift flag*: the SRL story and the judicial-decision-making material reach audiences owned by 06 and 05 respectively — coordinate in Phase B.

A6. Charge coverage. Subcommittee charge (how AI is/will transform legal work and judicial functions over ~24 months; capabilities): fully on capabilities, but over-inclusive on misuse relative to the explicit scope limit. Five objectives: innovation strongly served; access-to-justice served (SRL ladder); integrity/competence/confidentiality served only via the (out-of-scope) risk callouts.

A7. Citation anchors. [Transformational, § "High-Level Use Cases", "innumerable"]; [Transformational, § "Poor Legal Judgment", "75% answer"]; [Transformational, § "Self-Represented Litigant", "Tiered Reliability"]; [Transformational, § "Future Uses", "in the loop ... on the loop ... in command ... out of the loop"]; [Transformational, § "Institutional Problems", "institutional capture over the legal process"].

---

### 3. Generative AI Resources (04)

*Analyzed as a catalog, from the standpoint of the Resources subcommittee draft (Moriarty/Banks/Jones). Per the charge framework, this is a first-class substantive section, not an appendix.*

A1. Core conclusions / guidance (catalog of instructional content).

- How to evaluate a tool: hands-on testing on matters you know; focus on *precision* and *recall*; "Marketing demos are built to showcase wins; your testing should be built to reveal misses"; vary clean vs. noisy inputs, test consistency, test "I don't know" honesty; the "post-it test" (sort a week's tasks into AI-helps / AI-fails / unsure).
- Vendor due diligence: a question checklist (foundation model & version; grounding in primary sources; guardrails; training-on-inputs / ZDR; data access, residence, retention/deletion; SOC 2 / ISO 27001 / HIPAA; breach history & notification; indemnification; update cadence) and a "documents to scrutinize" list (ToS, Privacy Policy, DPA, security whitepaper, certifications, subprocessor list, model card, insurance).

- Confidentiality by tool tier: free/consumer = highest risk and inappropriate for client-confidential material under Colo. RPC 1.6; team/business/enterprise = first meaningful protections; legal-specific platforms; on-prem/local = strongest control.
- Tool categories: workflow/specialty tools; tools aimed at non-lawyers (UPL line "debated"); open-weights/local models (confidentiality upside; hardware/quality/security tradeoffs, malicious-model files).
- Building & maintaining competence: the three habits (use regularly; watch others' misuse / sanctions; revisit policies) and the "AI scout" practice for any organization larger than a solo.
- Quick-reference & population-specific resources: standing-order/court-rule trackers; key publications (*Colorado Lawyer* archive; ABA Formal Op. 512; ABA Task Force Year 2 Report); AI-literacy resources; watermarking/provenance/detection (C2PA, SynthID, NIST, commercial detectors with false-positive caveat); sample firm-policy templates; and tailored guidance for solo/small firms, LLPs (C.R.C.P. 207), judicial officers/court staff, and pro se litigants/public.

A2. Key supporting premises / evidence. ABA Model Rule 1.1 cmt. 8 anchors the competence rationale; case anchors (*Mata v. Avianca*; *In re Marriage of Haibt*; a D. Colo. matter) ground the hallucination-sanctions warning; statistics frame the risk ("~90%" of identified lawyer hallucination cases are solo/small-firm; pro se litigants "now make up the majority" of fabricated-citation cases; "two federal judges withdrew rulings in 2025").

A3. Accuracy & quality assessment. High practical quality; the evaluation discipline and vendor checklist are the document's most actionable assets. Durability is the dominant concern — this section is the most perishable in the corpus: named products (Clio Duo, MyCase IQ, Copilot, Prosēi, Cetient, Courtroom5), specific URLs, publication dates, and tracker lists will age fast; mark the *practice* (e.g., "ground citations in primary sources") as the keeper and the *named tool* as the replaceable illustration. Sourcing flags: the "~90%" and "majority of fabricated-citation cases" figures are attributed to "some trackers"; the "two federal judges" claim is specific and date-bound — all three need firm citations before publication. The non-endorsement disclaimer ("Inclusion ... is not an endorsement") is appropriately repeated.

A4. Priorities — what this paper most wants preserved (inference). *Inferred from emphasis*: (i) the evaluation methodology (hands-on, precision/recall, post-it test) and (ii) the vendor/confidentiality due-diligence framework as the transferable, evergreen core; (iii) competence as an *ongoing* practice (the AI-scout role). Textual basis: these are the most developed, instructional passages, contrasted with the lightly annotated link lists.

A5. Target audience. *Primary*: practicing legal professionals choosing and governing tools (solo/small-firm framing is the most developed). *Secondary*: LLPs (explicit [LLP] tagging), judicial officers/court staff, and pro se litigants/public (each given a tailored subsection). *Drift flag*: shifts among these audiences are mostly well-signaled by subsection headers — a model the other papers could follow.

A6. Charge coverage. Subcommittee charge (AI resources available): fully answered, with instructional content mapping to the charge's core "AI resources available" issue and the reference lists mapping to the *optional* annotated-bibliography task. Five objectives: competent service and confidentiality strongly served; innovation served; access-to-justice served (pro se/LLP subsections); integrity served via sanctions-awareness.

A7. Citation anchors. [04 Resources, § "How to Evaluate an AI Tool", "post-it test"]; [04, § "Hands-On Testing", "built to reveal misses"]; [04, § "Questions to Ask a Vendor", "zero data retention"]; [04, § "Confidentiality Across Tool Types", "usually the wrong choice"]; [04, § "Staying Current", "appoint an AI scout"]; [04, § "Solo and Small Firms", "approximately 90%"].

---

#### 4. Professional Ethics and the UPL Rules (05)

*Analyzed from the standpoint of the Ethics/UPL subcommittee draft (Martinez/Lupica/Siebecker/Ward).*

A1. Core conclusions / guidance. AI does not change a lawyer's duties; it routes them through existing rules. Organized rule-by-rule under the CRPC, the Code of Judicial Conduct, and UPL. Anchored to the January 2026 AI amendments — new Scope [20A], amended Comment [8] and new Comment [9] to Rule 1.1. Core holdings: competence requires *competent use* (and may eventually require *considering AI*), not avoidance (1.1); consider client disclosure (1.4); address AI's effect on reasonable fees (1.5); make "reasonable efforts" to protect confidences, avoiding public-tool input of client-identifying data (1.6); verify every citation/holding — AI filings with fabricated authority violate 3.1/3.3 and draw sanctions; delegation to AI does not reduce supervisory responsibility (5.1–5.3). For judges: independence/impartiality and bias vigilance (Canon 1), AI literacy and clerk oversight (Canon 2), ex parte risk from inputting case facts (2.9), prejudgment/appearance (2.10), and AI-related disqualification (2.11). For UPL: the practice-of-law definition and CRCP 228–240; Rule 5.5; a safe-harbor for lawyer-reviewed AI output; and the OARC Non-Prosecution Policy (Sept. 2025), a ≤3-year pilot deprioritizing UPL prosecution of supervised, technology-enabled tools (excluding juvenile/criminal/parental-rights matters).

A2. Key supporting premises / evidence. Verbatim rule text plus the official amendment language; the OARC policy as the load-bearing UPL development; nationwide sanctions practice (monetary sanctions, verification declarations, disciplinary referrals) as the candor anchor.

A3. Accuracy & quality assessment. The most authority-grounded and rule-accurate paper; quoting current Colorado rule text and the 2026 amendments is its strength. Flags: (a) Rules 7.1 (lawyer-advertising) and 8.4(g) (bias) are *named via Comment [9]* but not developed — partial coverage [05 Ethics/UPL, § "Rule 1.1. Competence", "communications concerning a lawyer's services ... and bias"]. (b) Perishable items to mark: the OARC pilot's three-year sunset and the "pilot" status itself — durable as a citation but explicitly time-limited [05, § "Colorado UPL Non-Prosecution Policy", "no more than three years"]. (c) Minor typo in source ("cannons" for "canons"). No accuracy errors of substance detected.

A4. Priorities — what this paper most wants preserved (inference). *Inferred from structure and reiteration*: (i) the rule-by-rule mapping as the document's compliance spine; (ii) the verify-every-citation duty (repeated across 1.1, 3.1/3.3, and the judicial track); (iii) the distinct judicial-officer track under the Code of Judicial Conduct; (iv) the OARC safe-harbor / Non-Prosecution Policy as the access-to-justice hinge of UPL. Textual basis: the rule-anchored organization itself and the recurring "Recommended Practices" boxes.

A5. Target audience. *Primary*: licensed Colorado lawyers (organized by the CRPC). *Distinct internal track*: judicial officers (Code of Judicial Conduct) — a self-contained audience within the paper. *Secondary*: developers/lawyers building or partnering with AI legal-service providers (Rule 5.5 / UPL). This matches the charge framework's note that the judicial-officer rule track *lives in 05*, not 07.

A6. Charge coverage. Subcommittee charge (CRPC, Code of Judicial Conduct, UPL): fully answered (with 7.1/8.4(g) only partially developed). Five objectives: integrity, competent service, and confidentiality strongly served; access-to-justice served via the UPL safe-harbor / OARC policy; innovation acknowledged (competence duty doesn't require avoidance).

A7. Citation anchors. [05, § "Scope [20A]", "does not diminish a lawyer's responsibilities"]; [05, § "Rule 1.1. Competence", "use them with knowledge, skill, and appropriate verification"]; [05, § "Rule 1.6. Confidentiality", "make reasonable efforts to prevent"]; [05, § "Rules 3.1 and 3.3", "fabricated case citations"]; [05, § "Rule 2.9: Ex Parte Communications", "case-specific facts"]; [05, § "Safe Harbor Provision", "reviewed, verified, and adopted"]; [05, § "Colorado UPL Non-Prosecution Policy", "deprioritize UPL prosecution"].

---

## 5. Self-Represented Litigants — Part 1 (06)

*Analyzed from the standpoint of the SRL subcommittee draft (Lupica/Dorancy/Riehl/Siebecker). This is the institution-, developer-, and court-facing half of a two-part set; its litigant-facing companion is 06a, analyzed at §6 below.*

A1. Core conclusions / guidance. Frames AI guidance for SRLs entirely around closing the access-to-justice gap, answering six questions (responsible development, accuracy, risks, steering to trustworthy tools, fit with existing infrastructure, avoiding a wider gap) plus what courts owe SRLs. Operative content: an ordered set of Guiding Principles (justice as overarching goal; integrity *and* access point the same way; informed agency over outcome-optimization; substantive over formal equality); Standards for Responsible Development (curated jurisdiction-specific knowledge bases, RAG with source pinning, automated citation verification, confidence/uncertainty flags, non-waivable fifth-grade-reading-level disclosure, mandatory pre-finalization review prompts, active referral pathways, mobile-first/multilingual/crisis-aware design, data minimization, incident-response, disaggregated performance metrics, community governance); a catalog of Risks (hallucination/procedural error, false confidence, systematic bias, opacity of failure, over-reliance/automation bias, data exposure, strategic error); Verification requirements; Steering Users to Trusted Tools ("a trusted tool that is harder to use than ChatGPT will lose"); Integrating Colorado's infrastructure (Sherlocks, LawHelp Colorado, self-help centers, clinics); Preventing a Two-Tier Justice System; Responsibility of the Courts (own the access problem; endorse responsible tools; collect/publish disaggregated outcome data; train judges/staff; fund what they endorse); and a flagged Privilege Question.

A2. Key supporting premises / evidence. Grounds claims in current Colorado practice ("not aspirational ... each guideline maps to a specific, auditable requirement"); design rationale rooted in the lived constraints of low-income, low-literacy, non-English, crisis-state users; specific failure modes (legal hallucination, jurisdiction bleed, temporal error, issue misclassification) tied to mitigations (bounded retrieval, source pinning, attorney sign-off).

A3. Accuracy & quality assessment. Internally rigorous and the most prescriptive paper (auditable "must" requirements). Flags: (a) Normative-conflict signal for Phase B — the claim "Restricting AI tools does not protect the integrity of the judicial system" [06 SRL, § "Integrity", "does not protect the integrity"] reframes integrity as aligned with access; this sits in tension with 05's integrity-first, restraint-oriented posture and should be surfaced on the conflict map, not resolved here. (b) Unfinished placeholders — bracketed "[designated body]" appears in the accountability requirements [06 SRL, § "Ongoing Accountability",

"[designated body]"]. (c) Some content is policy advocacy to institutions (court funding, liability frameworks, governance representation) rather than guidance to end-user audiences — within the subcommittee's "best practices for designing/developing tools" charge, but Phase B should confirm it belongs in a court-facing guidance document vs. a recommendations annex. (d) An apparent OCR/draft artifact: "court Sherlock staff" appears where "Sherlocks" (Colorado self-help staff) is meant — verify term [06 SRL, § "Leverage Trusted Intermediaries", "court Sherlock staff"].

A4. Priorities — what this paper most wants preserved (inference). *Inferred from the ordered-values framing and repetition*: (i) access-to-justice as the controlling value, with integrity reframed as compatible; (ii) design-for-the-least-advantaged-user (mobile-first, multilingual, fifth-grade reading level, crisis-aware); (iii) traceable, auditable obligations over aspiration; (iv) the unresolved privilege question as requiring formal Supreme Court/legislative action. Textual basis: the explicit "ordered set of values," the "auditable requirement" insistence, and the standalone Privilege section.

A5. Target audience. *Primary*: tool developers and the courts (the obligations run to builders and institutions), *on behalf of* self-represented litigants and the public as ultimate beneficiaries. *Secondary*: legal-aid organizations, Sherlock/self-help staff, and policymakers. *Drift flag*: the SRL is the *subject* but rarely the direct *addressee*; the document speaks about SRLs to those who build for and adjudicate over them — note this when locating audience labels in Phase B.

A6. Charge coverage. Subcommittee charge (A2J promise; minimize inaccurate/biased outputs for SRLs; best practices for designing/developing SRL tools; avoid widening the gap; quality-control standards for public-facing tools): fully answered, arguably over-inclusive on institutional policy. Five objectives: access-to-justice strongly (the paper's spine); integrity, competent service, confidentiality served via the development/verification/privacy standards; innovation served (it tells builders how to build responsibly).

A7. Citation anchors. [06 SRL, § "Justice as the Overarching Goal", "close the access-to-justice gap"]; [06 SRL, § "Integrity", "does not protect the integrity"]; [06 SRL, § "Transparency and Disclosure", "fifth grade reading level"]; [06 SRL, § "Harm Mitigation", "cannot hallucinate a citation that does not exist in its corpus"]; [06 SRL, § "Steering Users to Trusted Tools", "will lose to these platforms"]; [06 SRL, § "The Privilege Question", "unsettled question"].

6. Self-Represented Litigants — Part 2 (06a)

*Analyzed from the standpoint of the SRL subcommittee draft (Lupica/Dorancy/Riehl/Siebecker). The litigant-facing companion to §5.*

A1. Core conclusions / guidance. A plain-language, second-person guide — "Using AI to Help with Legal Problems: A Guide for People Who Represent Themselves" (draft 5/22/2026) — that translates Part 1's institutional standards into direct advice for the litigant. Operative content: (1) two very different kinds of tools — general-purpose platforms (ChatGPT, Google Gemini, Claude; "use with caution," not built for legal help, can hallucinate, no duty to flag ignorance, may store/train on your data) versus purpose-built legal tools (Colorado-specific, attorney-reviewed, user-tested — "what to look for"); (2) questions to ask before relying on a tool (Is a licensed attorney involved? Does it tell you it is not a lawyer? Does it state its area of law and last-updated date? Does it show its sources? What happens to your information? Does it tell you when it cannot help? Is it court/legal-aid endorsed?); (3) risks to watch for (AI makes things up; confident ≠ correct; it may not know what it does not know; it may not ask the right questions; procedurally correct ≠ strategically sound; your information may not be private; "you are least able to catch the mistakes you can least afford"); (4) extra steps before filing (read it carefully, verify every citation, check your deadline, get human help if stakes are high); (5) your privacy rights and the warning that AI-communication privilege is unsettled in Colorado — assume what you share is not confidential the way it would be with a lawyer; (6) where to get human help (LawHelp Colorado, Sherlock staff, Community Economic Defense Project, Colorado Legal Services, law school clinics, public libraries); (7) what courts should provide for you; and a five-point "Bottom Line" box.

A2. Key supporting premises / evidence. Mirrors Part 1's substance (identical two-tier classification, the same failure modes, the same disclosure/privacy/referral expectations) reframed as consumer literacy with concrete lay scenarios — e.g., a habitability defense a tenant might not know to raise, and the heuristic that an answer that "feels too clear and easy" is a reason to re-verify, not to stop. Authority is implicit rather than cited, which suits the audience.

A3. Accuracy & quality assessment. Substantively consistent with Part 1 and with 03/04/05; tone and reading level are well-calibrated — it *practices* the plain-language, fifth-grade-level disclosure standard that Part 1 only prescribes. Flags: (a) one Section 2 heading is still wrapped in editorial brackets — "[Is it endorsed by a court or legal aid organization?]" [06a Part 2, § "Section 2", "[Is it endorsed by a court]" — resolve before publication. (b) Named legal-aid organizations and a URL ([lawhelp.org/co](http://lawhelp.org/co); Community Economic Defense Project; Colorado Legal Services) are perishable specifics to date-stamp and maintain. (c) The privilege guidance ("assume ... not ... confidential") is the litigant-facing translation of

Part 1's open Privilege Question — accurate and appropriately cautious, but its correctness tracks the unresolved Supreme Court/legislature question; keep the two synchronized so they never diverge. (d) "Sherlock staff" recurs — confirm the intended Colorado term (carried over from Part 1).

A4. Priorities — what this paper most wants preserved. The "Bottom Line" box is the author's own explicit priority statement (so closer to *stated* than *inferred*): prefer purpose-built over general-purpose tools; require an up-front "not a lawyer" disclosure; look for attorney review and a last-updated date; never file without reading and verifying citations; know what happens to your information before sharing it. Reinforced by repetition across the risk and privacy sections.

A5. Target audience. *Primary*: self-represented litigants and members of the public, addressed directly in the second person — the corpus's only sustained public-facing register and the cleanest fit for the charge's "members of the public" audience. *Secondary*: the intermediaries who will hand it over (Sherlock/self-help staff, legal aid, librarians). *No drift* — the audience is consistent end to end, a model the rest of the corpus could emulate.

A6. Charge coverage. Subcommittee charge (A2J; minimize inaccurate/biased outputs reaching SRLs; quality signals for public-facing tools): fully answered from the consumer side. Five objectives: access-to-justice strongly; competent service and confidentiality served as user-protective literacy; integrity and innovation served indirectly. Phase-B note: 06 (Part 1, institution/developer/court-facing) and 06a (Part 2, litigant-facing) are the *same subject matter divided cleanly by audience* — the corpus's clearest worked example of the "resolve overlaps by audience first" principle. They are complementary, not duplicative, and should be kept as a paired set (a developer/court track plus a public-facing companion or appendix), with their shared substrate — two-tier classification, failure modes, disclosure/privacy/referral standards, and the privilege question — maintained in sync.

A7. Citation anchors. [06a Part 2, intro, "A Guide for People Who Represent Themselves"]; [06a Part 2, § "Section 1", "Use With Caution"]; [06a Part 2, § "Section 2", "Is a licensed attorney involved?"]; [06a Part 2, § "Section 3", "Confident does not mean correct"]; [06a Part 2, § "Section 5", "not yet settled in Colorado"]; [06a Part 2, § "Section 6", "LawHelp Colorado"]; [06a Part 2, § "Quick Reference", "The Bottom Line"].

---

## 7. Generative AI and Legal Professionals (07)

*Analyzed from the standpoint of the Legal Professionals subcommittee draft (Ward/Banks/Dorancy/Jones/Riehl).*

A1. Core conclusions / guidance. Operational best practices across ten sections: (1) introduction/scope/foundational principles (AI is probabilistic, "does not 'know' the law"); (2) technological competence (open-loop vs. closed-loop systems; understand tokenization/hallucination); (3) confidentiality (thorough anonymization if a public tool must be used; prefer enterprise-grade ZDR/no-training/encryption; AI-tailored incident-response); (4) verification, accuracy, and duty of candor ("AI should be treated as an assistant, never as an ultimate, unquestionable authority"; independent human verification of every citation/quote/fact); (5) supervision (firm AI policy defining permitted data categories and mandatory verification; supervisors must ask whether AI was used; training culture); (6) client disclosure/consent/communication (discuss at outset; informed consent before uploading sensitive volumes; engagement-letter clause); (7) billing/fees/economics (don't bill saved hours; *do* bill prompt-crafting and review; subscriptions are overhead; case-specific tools may pass through with consent; alternative fee arrangements); (8) bias/fairness/access to justice (review outputs for algorithmic bias; deploy AI to close the A2J gap); (9) judicial officers & court administration (adjudicative core is "exclusively human"; administrative uses permitted with verification; standing orders; clerk data-security); (10) attorneys/LLLPs/paralegals (front-line adoption; coordinated team responsibility).

A2. Key supporting premises / evidence. The probabilistic-text premise drives the verification mandate; the anonymization and enterprise-procurement guidance operationalizes confidentiality; the billing section reasons from "fees must be reasonable and accurately reflect the actual work performed"; judicial guidance reasons from independence/impartiality and record security.

A3. Accuracy & quality assessment. Practically comprehensive and readable. Flags: (a) Heavy overlap with 05 on confidentiality, verification/candor, supervision, fees, and bias — 07 is operational where 05 is rule-anchored; this is the central consolidate-vs-divide-by-aspect decision for Phase B. (b) Audience drift — titled "by Legal Professionals" but Section 1 addresses "judicial officers, attorneys, LLLPs, paralegals, and court staff," and Section 9 is a full judicial-officer/court-administration track that overlaps 05's Code-of-Judicial-Conduct track [07, § "Section 9", "Judicial Officers and Court Administration"]. (c) Technical simplification — the "open-loop vs. closed-loop" dichotomy is a serviceable but imprecise gloss relative to 03's more accurate account; harmonize terminology with Fundamentals. (d) Perishable: enterprise-tool/security-certification specifics. No fabricated authority detected.

A4. Priorities — what this paper most wants preserved (inference). *Inferred from repetition and detail*: (i) independent human verification as the non-negotiable core (echoed across

§§2, 4, 5, 9); (ii) the firm-level AI policy as the governance mechanism; (iii) the detailed billing/economics treatment (its most distinctive, developed contribution); (iv) the "adjudicative core is exclusively human" line for judges. Textual basis: §4's "most critical advisory mandate" language and the unusually granular billing section.

A5. Target audience. *Primary*: practicing lawyers, plus LLLPs and paralegals (named in §10) and supervisors. *Distinct track*: judicial officers and court administrators/clerks (§9). *Drift flag*: per the charge framework, 07 "reaches beyond lawyers to judicial users and more general practice-of-law concerns" — confirmed; the judicial material here overlaps 05's judicial track and must be reconciled by audience in Phase B (the framework's resolve-by-audience-first rule).

A6. Charge coverage. Subcommittee charge (practical guidance/best practices for responsible integration into profession and judiciary; effect on competencies and ethics; efficiency/cost/A2J); risks; operational standards and model policies): fully answered. Five objectives: competent service, integrity, and confidentiality strongly served operationally; access-to-justice served (§8); innovation served (efficiency/billing). The "model policies" element of the charge is gestured at (a firm AI policy is described) but no full sample policy text was retrieved — confirm whether one exists; if not, flag as a partial item that the charge's optional "sample AI policy" task could fill.

A7. Citation anchors. [07, § "Section 1", "does not 'know' the law"]; [07, § "Section 2", "open-loop ... closed-loop systems"]; [07, § "Section 3", "comprehensive anonymization practices"]; [07, § "Section 4", "never as an ultimate, unquestionable authority"]; [07, § "Section 7", "should only be billed for those two hours"]; [07, § "Section 9", "exclusively human endeavors"].

---

Phase-A-level cross-references for Phase B (recorded, not resolved)

*These were surfaced descriptively during Phase A. Their resolutions now live in the separate Phase B Design Document (section order, the Risks phenomenon/obligation boundary, the 06+06a merge direction, the UPL information-vs-advocacy split, the two-question resolution of the caution/access tension, the single Judiciary home, and the evergreen/annex convention). They are retained here unchanged as the descriptive record of what the corpus presented.*

- 05 ↔ 07 overlap confirmed on confidentiality, verification/candor, supervision, fees, bias; 05 rule-anchored, 07 operational — *divide by aspect/audience per the framework's resolve-by-audience rule.*

- 05 (Code of Judicial Conduct) ↔ 07 §9 — both run a judicial-officer track; reconcile ownership.
- Risk/misuse has no clean home: scattered across 03 (architectural why), Transformational (out-of-charge inline callouts + Institutional Problems), 05 (rule violation), 06 (SRL-specific risks), 07 (operational). A dedicated risk treatment may be warranted — flagged, not forced.
- Value tension: 06's "restriction does not protect integrity" vs. 05's integrity-first restraint — conflict-map item.
- Evergreen exposure (descending): 04 (highest — named tools/links/stats) > Transformational (model behaviors, dollar figures) > 03 (token/parameter specifics) > 07 (enterprise specifics) > 05 / 06 (most durable; perishables limited to the OARC pilot sunset and bracketed placeholders).
- Sourcing to firm up before publication: 04's "~90%/" "majority"/"two federal judges 2025"; Transformational's "more accurate than human judges."

## Phase B Design Document

### Colorado LTAC AI Guidance Document — setup brief for the synthesis & drafting session

#### 0. How to use this document

This is the controlling brief for **Phase B**, the session that turns the seven analyzed drafts into one unified guidance document. It is paired with the **Phase A Analyses** artifact, which is the descriptive record of the source corpus — use Phase A to know what each source actually says, what each section most wants preserved (its A4 priorities), and which citation anchors (its A7 set) to carry forward so drafted text stays traceable to its source. This document records the decisions already settled with the Committee lead so the drafting session does not relitigate them, states the global standards the draft must meet, and lists what Phase B must still produce and what is reserved for the full Committee.

---

#### 1. Controlling authority and scoring rubric

Controlling authority remains the **Chief Justice's charge** (doc 01). Every structural and editorial choice is scored against the charge's five objectives, which are the rubric throughout: safeguard the integrity of the legal system; promote competent client service; protect client confidences; support access to justice; and encourage innovation — consistent with the Colorado Rules of Professional Conduct. Where a section's preserved priority (Phase A's A4) collides with these objectives or with accuracy (A3), the objectives win and the reason is stated, not left implicit.

**Audience set** (from the charge): legal professionals, judicial officers, self-represented litigants, and members of the public.

**Two reading modes, served at once.** *Standalone* — a reader who jumps to one section finds it self-contained. *Cover-to-cover* — sections build pedagogically, capability before consequence. The architecture in §3 is designed so each section rests on the substrate the earlier sections lay down; that is what lets a section stay shallow on prior material without heavy cross-referencing.

---

#### 2. Non-negotiable source hygiene (carry into drafting)

These survive from the original framework and bind the drafting session:

- **Cite content, not filenames.** Anchor each citation to (a) source short-title, (b) nearest heading, (c) a short distinctive phrase. No fabricated page or line numbers.
  - **No invented citations or authorities.** If uncertain, omit.
  - **No manufactured deliberation.** Write from a standpoint, labeled as such; never stage a fictional debate.
  - **Keep ground truth separate from outside research.** Anything from outside the corpus goes in a marked Supplementary block with its own citation, never attributed to a subcommittee.
  - **Flag gaps; do not paper over them.**
- 

### 3. Settled architecture

The document moves capability → consequence → application → audience-specific duties, then quarantines what ages fastest. Each section's *owns / draws from / boundary* is fixed; final sequencing *within* the rules cluster is a Phase B outline task (§6).

**1. Foundations** (*from 03 Fundamentals*) — what the technology is and the architectural "why" that makes later risks and benefits intelligible. Owns: the conceptual model of an LLM, its capabilities, and its built-in limitations. Boundary: explains *why* failure is possible; does not catalog legal failures (that is Risks) or rules (that is the rules cluster).

**2. Transformations** (*from 02*) — the optimism register. How the technology is used today and will likely be used over the next ~24 months, carried by the user stories. Owns: capability classes and the imaginative case for adoption. Boundary: keeps only the **benefit** halves of its current Benefit/Risk callouts; the cautionary halves migrate to Risks. Ends on the one invariant that hands off to the next section — capability and risk grow together.

**3. Risks** (*NEW — extracted from 02's risk callouts, grounded in 03's limitations*) — the cold water. Reorganized **by risk type** (e.g., hallucination, fabricated/deepfake evidence, data leakage/confidentiality loss, automation bias / over-reliance, surrendering judgment, single-shot error), each tied back to the architectural cause in Foundations and paired with a general mitigation. Owns: the **phenomenon** layer (what the failure is, how it shows up concretely in legal work, and the audience-independent mitigation). Boundary: see §4.1 — it must *not* re-absorb rule analysis.

**4. Resources** (*from 04*) — how to put the mitigations into practice: evaluating, choosing, configuring, and governing tools; building and maintaining competence. Owns: the evaluation methodology (hands-on testing, precision/recall, the post-it test),

vendor/confidentiality due-diligence, and the competence practices. Depends on Risks (it applies exactly the mitigations Risks defines), which is one reason Risks must precede it.

**5. Rules-and-audience cluster** — each subsection pulls the Foundations→Transformations→Risks→Resources threads forward to a specific audience or rule set, going only as deep as that audience needs:

- **Professional Ethics & the Rules** (*from 05; the lawyer-facing rule analysis*) — the CRPC track (1.1, 1.4, 1.5, 1.6, 3.1/3.3, 5.1–5.3, plus 7.1/8.4(g)), the 2026 AI amendments, and the **lawyer-facing UPL framework** (Rule 5.5, the practice-of-law definition, the lawyer-supervision safe harbor).
- **The Judiciary** (*single home for judicial conduct — consolidating 05's Code of Judicial Conduct track and 07 §9's operational judicial material*) — see §4.6.
- **Legal Professionals in Practice** (*from 07, operational*) — the operational "how" for lawyers, LLLPs, and paralegals: supervision, client disclosure, anonymization, billing/economics. Divided from the Ethics section by aspect (rule/why vs. operational/how) per §4.7.
- **Self-Represented Litigants** (*merged 06 + 06a; see §4.3*) — a pro se user guide, ending with the access-relevant UPL facts and the balancing discussion in §4.5.

**6. Annex — Current Resources and Examples (as of [date]; reviewed [cadence])** — the quarantine for perishables (§5.3).

---

## 4. Key editorial decisions

### 4.1 Risks: extraction and the phenomenon/obligation boundary

The risk material now in Transformations is not a clean block — it is the "Risk:" halves of the Benefit/Risk callouts woven through the user stories. Extracting it means **reorganizing by risk type**, not lifting a section. The dividing line that keeps Risks from becoming either a toothless list or a duplicate of the rules:

- **Risks owns the audience-independent layer:** what the failure mode is, how it manifests in legal work, and the general mitigation. A hallucination behaves the same whether a lawyer, judge, or pro se litigant hits it; mechanism and mitigation generalize.
- **The rules cluster owns the audience-specific layer:** how that same failure becomes a Rule 3.3 candor problem, a Rule 1.6 confidentiality problem, or a

judicial-impartiality problem, and what the governing rule then demands. Discipline: Risks states phenomenon and mitigation; obligation and consequence live downstream. If Risks starts reciting rules, it recreates the scope creep that put this material in Transformations in the first place.

#### **4.2 Risks is standalone and sits *before* the rules cluster (resolved)**

The open consideration — fold Risks into the rule sections, or keep it standalone first — is resolved in favor of **standalone, before**. It gives risk one authoritative home (fixing the scatter Phase A flagged), satisfies the dependency that Resources applies the mitigations Risks defines, and honors the cross-reference restraint (§4.8): the rule sections build a normative layer on a substrate Risks already laid, rather than pointing back to it.

#### **4.3 Merge of 06 + 06a — direction matters**

The unified self-represented-litigant section is built on **06a's litigant-facing register as the spine** (an expanded Part 2), into which the *keepable factual core* of Part 1 is pulled: the two-tier tool classification (general-purpose vs. purpose-built), the failure-mode taxonomy, verify-before-filing, privacy/privilege, and where-to-get-human-help. **Dropped:** Part 1's developer mandates and its rule-reform argument. The result should read like an expanded user guide, not a softened policy paper. Anchoring on Part 1's structure and bolting Part 2 on would fight the prescriptive voice the whole way; anchoring on Part 2 avoids that.

#### **4.4 UPL — preserve the information, remove the advocacy**

Split the UPL material along the audience seam:

- **Lawyer-facing framework** (Rule 5.5, the practice-of-law definition, the lawyer-supervision safe harbor, the duty not to assist UPL when building/partnering) stays in **Professional Ethics & the Rules** as obligation-layer content. Nothing factual leaves it.
- **Access-relevant facts** for pro se users and builders — chiefly that the **OARC non-prosecution policy** currently deprioritizes enforcement against certain attorney-supervised, technology-enabled tools — go at the **end of the merged pro se section**, stated as the current regulatory state of play, not as a brief to change it.
- **Transformations** may carry only a purely factual forward note: the OARC policy is a time-limited pilot under evaluation, so the permissible envelope may shift. Fact, not demand. Do not relocate the broader UPL framework into Transformations; it tends to drift back toward advocacy there.

#### **4.5 The two questions, the tradeoff hinge, and the balancing factors**

The 05/06 disagreement is not one question with two answers; it is **two questions, each answered correctly in its own track**:

- *Rules/ethics track (lawyer- and judge-facing)*: given a licensed professional who bears responsibility and is equipped to catch errors, how should they deploy AI consistent with their duties? Answer: restraint-first — caution is genuinely protective because a trained human-in-the-loop absorbs the residual risk.
- *Pro se track (public-facing)*: given a person with no professional backstop and the most to lose, how can AI be used so it widens rather than narrows access? Answer: blanket restraint protects no one and entrenches the gap; the work is in tool design, lay-appropriate verification, and knowing when to stop and get a human.

The hinge where the tracks meet is best phrased as a **tradeoff, presented as a balance rather than a verdict**: *restricting the proper and informed use of AI trades access for integrity*. Its natural home is the pro se / UPL discussion, which should lay out the corpus-grounded **factors that tip the balance** and let the reader (and the Committee) see the structure:

- whether a licensed lawyer reviews the output (the safe-harbor axis);
- the stakes and reversibility of the matter (the corpus's "high stakes → get human help");
- whether the tool is purpose-built with bounded retrieval or a general-purpose model;
- the user's realistic ability to verify;
- the area of law — the OARC policy's own carve-outs (juvenile, criminal, termination of parental rights) are a corpus-grounded signal of where the balance already tips toward caution.

Reconcile or cut the bald line "restricting AI does not protect the integrity of the judicial system." In a user-facing register it becomes: *used carelessly, AI can hurt your case; used well, it can help — here is what moves the needle*.

#### **4.6 Judicial conduct: one home**

Judicial conduct is discussed **once**. Consolidate 05's Code of Judicial Conduct track (Canons 1–2; Rules 2.9 ex parte, 2.10 pending matters, 2.11 disqualification; clerk supervision) and 07 §9's operational judicial material (administrative-use boundaries, "adjudicative core is exclusively human," standing orders, clerk data security) into a single

**Judiciary** home. Default placement is within the rules cluster, since the Code of Judicial Conduct is the rule-anchored spine and the operational material attaches to it.

#### **4.7 05 ↔ 07 overlap: divide by audience/aspect**

Confirmed resolution: **divide, do not consolidate**. 05 is rule-anchored (the *why*, for lawyers and judges); 07 is operational (the *how*). The shared topics — confidentiality, verification/candor, supervision, fees, bias — split cleanly by aspect, with the judicial slice routed to the single Judiciary home per §4.6.

#### **4.8 Cross-reference restraint**

Avoid a web of explicit "see Section X" pointers. Each section goes as deep as its own subject requires and relies on the substrate earlier sections established. The architecture is built to make this possible; lean on sequence, not signposting.

---

### **5. Global standards for the drafted document**

#### **5.1 Reading level — approximately high-school-freshman (≈9th grade)**

The body must be **accurate but accessible**, written so a ninth-grade reader can follow it. This is the document's own target. Note one thing this is *not*: the "fifth-grade reading level" standard the corpus recommends is a requirement Phase B should keep recommending **for AI tools' own user disclosures**, addressed to the most vulnerable users — it is guidance *about tools*, not the reading level of this document. Keep the distinction explicit so a future editor does not collapse the two.

#### **5.2 Unified voice**

The source drafts arrive in very different voices — technical-metaphorical (Foundations), narrative (Transformations), checklist (Resources), rule-recitation (Ethics), prescriptive (06), plain second-person (06a), formal advisory (07). The drafted document should read as **one author** to the extent possible: plain, direct, confident, and explain-then-illustrate. 06a is the closest existing model for the target tone. Bring the rule recitations and the technical metaphors toward that register **without sacrificing accuracy** — state the durable, precise point plainly in the body and let footnotes carry the technical or citation detail (§5.3).

#### **5.3 Evergreen / annex convention**

The principle is **quarantine, not deletion**: every perishable fact stays, but moves to a slot built to be swapped, so the surrounding prose goes near-evergreen. Four moves:

1. **One dated annex.** Route the exhaustive perishables into "Current Resources and Examples (as of [date])": named products; training/publication lists; current sanctions cases; raw statistics used illustratively; and time-stamped regulatory edges (e.g., the OARC pilot's sunset window). The annex holds *instantiations*; **load-bearing legal facts** (that the safe harbor exists, that the OARC policy exists, that the 2026 RPC amendments exist) stay in the body.
2. **Split keeper from illustration at the sentence level.** State the durable principle in body prose; put its current instantiation in a dated inline box or footnote. Examples: *Resources* — body says "prefer tools that ground citations in primary sources," a dated box names the products that currently do; *Foundations* — body says "models consider only a bounded amount of text and degrade before that limit," a footnote carries the ~1M-token figure; *Ethics* — body states the permanent verify-every-citation duty, a footnote carries the dated 58–88% hallucination study.
3. **Reframe Transformations from snapshot to trajectory.** Keep the user stories (they teach a workflow pattern — check ground truth, treat output as a ~75% draft you own the last 25% of) but genericize the model-specific numbers and illustrative figures ("\$300," "ninety minutes") as illustrative-at-time-of-writing; speak in capability classes plus the capability-and-risk-grow-together invariant.
4. **Anchor to durable substrate.** Where you can choose, hang a point on a rule, a duty, or an architectural fact from Foundations rather than a case or a product — that is why the rules track ages best. Cite the perishable in a footnote; state the permanent thing in the body.

**Maintenance, built in.** The charge already requires recommendations at least annually; tie the annex's review cadence to that cycle and include a short maintainer's note (what to check each cycle, where it lives) so the next editor updates one annex and a few dated boxes rather than re-reading the whole document. **Guardrail:** do not over-quarantine — keep one inline illustration per concept so the standalone reader is not sent to an appendix for a usable example; send only the exhaustive, maintained lists downstream.

#### 5.4 Dual reading modes

Restated as a drafting test because it interacts with everything above: every section must work for a reader who lands on it cold, *and* in sequence. Self-containment plus the no-heavy-cross-reference rule means a section may briefly restate a prior point in its own words where comprehension demands it — that is not duplication to be stripped.

---

## 6. What Phase B must still produce

Per the original framework, Phase B delivers four products. Several are now partly pre-settled; finish them and record the reasoning.

1. **Overlap map** — largely resolved (05↔07 divide by audience §4.7; judicial single-home §4.6; 06↔06a merge §4.3; risk consolidation §4.1–4.2). Confirm no new overlaps surfaced and write the map.
2. **Conflict map** — the principal conflict (the 05/06 caution-vs-access tension) is articulated and resolved by audience-split + the balancing framing in §4.5. Scan for any remaining conflicts (e.g., the open/closed-loop simplification in 07 vs. Foundations' more precise account) and record them.
3. **Resolution criteria, in order:** (i) fidelity to charge + five objectives; (ii) accuracy/quality from A3; (iii) avoiding duplication/reader confusion. Honor each A4 priority unless it loses on these — and say why if it loses.
4. **Annotated outline** — over-inclusive, heavily anchored (reuse the A7 anchors), losing nothing silently; simplification is a later pass. Specifically, the outline must finalize: the **sequence within the rules cluster**; whether **The Judiciary** is its own section or a track inside Ethics; placement of edge content (the privilege question; the institutional/court-responsibility material being demoted from 06; any sample policy or checklist); and which perishables go to the annex.

---

## 7. Open questions reserved for the full Committee

- **Form and format (charge-optional tasks):** whether to include an executive summary, checklists/forms, and a sample AI-use policy for practitioners (07 describes but does not supply a full sample policy; this could fill the charge's optional "sample policy" task), and where each would live.
- **The privilege question:** the corpus flags as unsettled whether a self-represented party's communications with an AI tool are privileged, and recommends the Supreme Court or legislature resolve it. Decide how the unified document presents this — as a neutral notice to readers plus a flagged recommendation — without itself adopting a position.
- **Court-responsibility / policy material demoted from 06:** whether any of it (court endorsement of tools, outcome-data collection, funding) belongs in a separate recommendations annex addressed to the Court rather than in the user-facing body.

- **Annex ownership and cadence:** who maintains "Current Resources and Examples," and confirmation that its review cycle is pinned to the annual recommendation deadline.
  - **Reading-level verification:** how/whether to test the final draft against the ≈9th-grade target.
- 

## 8. Using the Phase A artifact during drafting

Phase A stays the **descriptive record of the corpus** and is the drafting session's reference for three things: (1) **what to preserve** — each paper's A4 priorities mark what that source most wants kept; honor them unless they lose under §6's criteria. (2) **What to watch** — each A3 flags accuracy issues, unfinished placeholders (e.g., "[designated body]", the bracketed 06a heading), perishable specifics, and sourcing to firm up before publication (the 04 statistics, the Transformations "more accurate than human judges" claim, the leaked-note 03 footnote). (3) **How to stay traceable** — each A7 supplies the heading+phrase anchors to reuse so drafted prose links back to its source. Do not edit Phase A's analyses to match Phase B decisions; the two artifacts play different roles — Phase A says what the sources *are*, this document says what we will *build* from them.

## Change Catalog: Synthesized Guidance Document vs. Subcommittee Source Drafts

Sources compared: 01/01a (Establishing Document and Minutes), 02 (Transformational Aspects), 03 (Fundamentals), 04 (Resources), 05 (Ethics/UPL/Judiciary), 06 (SRL Part 1 — policy framework), 06a (SRL Part 2 — consumer guide), 07 (Legal Professionals).

---

### 1. How the Sections Were Synthesized and Rearranged

#### 1.1 — Overall architecture: eight drafts collapsed into six Parts with assigned roles

- The standalone subcommittee drafts were converted into a single document where each Part is given an explicit, non-overlapping job, announced in connective text that exists in no source: Part 1 explains *why* failure is possible; Part 2 shows capability; Part 3 catalogs risk-with-mitigation; Part 4 supplies practical method; Part 5 maps duties to rules; Part 6 speaks to the public. Cross-references ("taken up in the Parts that follow," "the working posture illustrated in the previous Part") stitch them together.
- **Alignment:** The minutes contemplated exactly this — separate drafting subcommittees "but ... a single final editor," and Phase II of the action plan called for review "for consistency and clarity." The Part structure also tracks the "Core Issues" list in the establishing document almost one-for-one (fundamentals → transformation → resources → rules implications → misuse/risk).

#### 1.2 — Document 07 (Legal Professionals) was dissolved entirely and distributed

07 overlapped nearly everything else; the synthesis eliminated it as a freestanding unit and merged its sections into the Parts where the same subject already lived:

- §§1–2 (probabilistic nature, tokenization, RAG, competence) → absorbed into Part 1 and Part 4 §3(c) (RAG explanation now sits in "Legal-Specific Platforms").
- §3 (anonymization, enterprise tools, incident response) → merged into the Rule 1.6 unit (Part 5), which now carries 07's specifics: "removing names, specific dates, account numbers, unusual geographic locations, and distinctive facts" and the new incident-response-plan bullet; the general data-exposure discussion went to Part 3 §4.
- §4 (verification and candor) → merged into Part 3 §1 (Hallucination) and the Rules 3.1/3.3 unit.

- §5 (supervision, written policies, "ask whether AI was used") → merged into the Rules 5.1/5.2/5.3 unit, which gained the recommended practice "affirmatively ask whether any AI tool was used."
- §6 (client disclosure, engagement-letter clause) → merged into the Rule 1.4 unit, which now carries 07's engagement-letter and "don't bury the client in jargon" language.
- §7 (billing, cost pass-through, flat fees) → merged into the Rule 1.5 unit (case-specific AI costs vs. overhead; alternative fee arrangements).
- §8 (bias and access to justice) → split between Part 3 §3 (Bias) and the Rule 8.4(g) unit.
- §9 (judicial officers and court administration) → appended to Part 5 §3 as the new subsection "Administrative Use and the Adjudicative Core," nearly verbatim (adjudicative core, standing orders, ordering production of suspect citations, clerk data-security, Sherlock-adjacent staff caution).
- §10 (attorneys/LLPs/paralegals) → its substance reappears in the new LLP unit (Part 5 §4) and in Part 3 mitigations; note the synthesis also standardizes 07's "LLLP" to Colorado's "LLP."
- **Alignment:** The Legal Professionals subcommittee's charge was to "translate technical realities into operational standards [and] best practices" — the synthesis treats that as cross-cutting material rather than a Part, which resolves the largest duplication problem among the drafts.

### 1.3 — Risks were extracted from everywhere and centralized in a new Part 3

- 02's trailing analytic sections (volume of low-quality output; evidentiary issues; "Institutional Problems") were lifted out of the Transformations draft and became Part 3 §§8, 6, and 9 respectively.
- 03's §4 ("Limitations of Proper Use" — hallucination, deception, bias, sycophancy) was split: the *causal* explanation stays in Part 1 §4, while the *legal-practice consequences and mitigations* were rewritten into Part 3 §§1–3, each now ending with a paired "general mitigation."
- 07's verification, confidentiality, and supervision risk discussions and 06's SRL risk list (false confidence, strategic error, autonomy illusion, verification gap) were folded into Part 3 §§4, 7 and echoed in plain language in Part 6 §3.

- New risk taxonomy headers were imposed (Hallucination; Poor Legal Judgment; Bias; Confidentiality; Agent Error; Fabricated Evidence; Over-Reliance; Volume; Institutional Capture), each tied "to its cause and paired with a general mitigation."
- **Alignment:** Directly serves the minutes' instruction that the Transformations subcommittee "focus on the capabilities of the technology rather than on how it has been used and misused" — misuse content that had crept into O2 was moved out — and serves the charge's core-issue item on "how legal professionals ... have misused AI."

#### 1.4 — User-story callouts were systematically relabeled and rebalanced

- In O2, the in-story callout boxes were mostly labeled **Risk** (e.g., "Risk: Confidentiality and Privilege Waiver," "Risk: Garbage In, Garbage Out," "Risk: Poor Legal Judgment," "Risk: Surrendering Judgment," "Risk: Fake Evidence," "Risk: Single-Shot Errors," "Risk: Delegation of Judicial Decision-Making"). In the synthesis nearly all were converted to **Best Practice** or **Benefit** boxes ("Best Practice: Protecting Confidentiality," "Best Practice: Skepticism," "Best Practice: No Surrender of Judgment," "Best Practice: Careful Prompt Creation," "Benefit: A Mirror, on Request," "Benefit: A Second Opinion Costs Almost Nothing"), with the underlying risk discussion relocated to Part 3 (e.g., "Surrendering Judgment" became Part 3 §7 and an in-narrative sentence: "Agents simulate; litigators decide"; "Fake Evidence" became Part 3 §6; "Single-Shot Errors" became the cross-check mitigation in Part 3 §1; "Delegation of Judicial Decision-Making" became the drafting paragraph in Part 5 §3).
- Smaller conflict-resolutions inside the stories: "damning thread" → "damaging thread"; the privilege-waiver claim was softened and made consistent with Parts 3, 5, and 6 by repeating the "currently unsettled question" framing rather than asserting waiver.
- **Alignment:** Keeps Part 2 within its capabilities-focused mandate while preserving the practical lessons; the consistent "unsettled" framing reflects the synthesis goal of objective information over conclusions of law the Committee has no authority to make.

#### 1.5 — Use-case list reorganized (Part 2 §1)

- O2's bulleted high-level use cases became numbered subsections 1(a)–1(n); "Analyzing Document Sets" was moved from second position to 1(e) so that Drafting and Research lead; the Research bullet's warning about early fabricated-citation

notoriety was trimmed to a neutral sentence about specialized vs. general tools (the fuller treatment now lives in Part 3 §1).

- A new framing note was inserted before the stories: dollar amounts, document counts, and timings "are illustrative as of the time of writing and will age," and "the durable point is the workflow pattern."

### **1.6 — Fundamentals (03) → Part 1, restructured and de-quantified**

- 03's "Error States" chapter was split into Part 1 §3 ("Built-In Limitations") and a new §4 ("Models can produce false or misleading content"), grouping hallucination, bias/prompt sensitivity, and deception/false reasoning as a family.
- The manipulation/jailbreaking material and the agentic-misbehavior material were separated into §3(c) and a new §3(d) ("Model misbehavior and misalignment"), which merges 03's footnoted examples (verbatim training-data extraction; the email-deleting agent) into body text.
- Perishable specifics were generalized: "on the order of a thousand neurons ... a hundred layers deep" → "many ... many deep"; "around one million tokens" → "large, but not large enough"; the blackmail steering statistic "from roughly twenty-two percent to roughly seventy-two percent" → "more than tripled the rate."
- 03's closing "nature of an LLM's existence" passage became Part 1 §5 with a new bridge ("These are some of the reasons why a human must stay in the loop...") connecting it forward to Parts 3–5.
- **Alignment:** Matches the Fundamentals subcommittee charge (definition, technical foundations, capabilities and limitations) while serving the evergreen-document concern reflected in the minutes' observation that other jurisdictions' documents age quickly.

### **1.7 — Resources (04) → Part 4, with the Annex device**

- 04's structure (evaluation → confidentiality tiers → tool categories → competence → populations) was retained nearly intact, but every named product, tracker, publication, training program, and template was moved to a referenced "Annex (Current Resources and Examples) ... refreshed on a set cycle," with the durable practice retained in the body.
- 04's pro-se evaluation paragraph was expanded into §1(d) ("The Increased Burden on Pro Se Users") and cross-referenced to the new Part 6.

- Statistics were softened to survive aging: "approximately 90% of identified ... hallucination cases" → "the large majority"; pro se litigants "make up the majority of cases involving fabricated citations" → "a growing share"; "two federal judges withdrew rulings in 2025" → "there have been instances of judges withdrawing rulings."
- **Alignment:** Preserves the Resources subcommittee's gathering work and the action plan's optional "annotated bibliography" task while honoring the Committee's repeated "inclusion is not an endorsement" caveat — important to the charge's integrity-of-the-system objective, since a court document naming commercial vendors in its body risks the appearance of endorsement.

### 1.8 — Ethics/UPL/Judiciary (05) → Part 5 §§1–3, augmented

- 05's three blocks (CRPC, Judiciary, UPL) were kept but reordered (CRPC → UPL → Judiciary → new LLP unit) and the rules were recast into the boxed rule-text + commentary + Recommended Practices format throughout.
- The Comment [9] text was made the organizing spine: "Each of the Rules named in Comment [9] is discussed below," which drove the addition of Rule 7.1 and Rule 8.4(g) units that the rule-by-rule treatment in 05 did not develop.
- The UPL section now leads with C.R.C.P. 232.2's definition — including subsection (c)(9) on interactive technology — before Rule 5.5, and the safe-harbor/non-prosecution discussion was generalized (see removals below) and merged with 05's dual "OR" recommended-practice lists into one coherent list.
- The judiciary unit absorbed 07 §9 as "Administrative Use and the Adjudicative Core," and the drafting-sharpens-decision-making caution from 02's Judge-story callout was rewritten into a dedicated paragraph there.
- **Alignment:** Serves the Ethics subcommittee's charge (CRPC, Code of Judicial Conduct, UPL) and the synthesis goal of completeness — discussing every rule Comment [9] names closes an obvious gap a reader of the January 2026 AI amendments would notice.

### 1.9 — The two SRL drafts (06, 06a) → Part 6, voice unified

- 06a (the plain-language consumer guide) supplied the skeleton and most of the text of Part 6 (two tool types; questions to ask; risks; resources).
- 06 (the normative policy framework) was not carried over as policy; instead its substance was *translated into consumer-facing signals*: the mandatory fifth-grade-

reading-level disclosure regime became "this warning should be simple enough that a fifth-grader could understand it ... if this disclosure is hidden or missing, that is a red flag" (Part 6) and "a signal of a responsibly built tool" (Part 4 §3(d)); the steering/autonomy standards became risk 3(e) ("The tool may push you toward one choice") and the bottom-line item "A good tool shows you options; it does not choose for you"; the verification-gap analysis became 3(h); the institutional-endorsement strategy became the "Is it endorsed by a court or legal aid organization?" question and the closing description of what courts are doing.

- **Alignment:** This is the clearest execution of the user's stated synthesis goal — "objective information instead of policy recommendations." It also honors the SRL subcommittee's access-to-justice charge while staying inside the Committee's actual mandate (guidance to users, not regulation of developers).

---

## 2. Content Removed Entirely

- **All named commercial tools, vendors, URLs, publications, trackers, and templates** (04's tool tables — ChatGPT, Claude, Gemini, Harvey, Westlaw CoCounsel, Lexis+ AI, Clio Duo, Prosēi, Courtroom5, Hello Divorce, Ollama, etc.; the *Colorado Lawyer* article list; Charlottin trackers; Ropes & Gray tracker; bar policy templates; Anthropic/OpenAI/Google/Microsoft training links; case-name list including *Crabill, Al-Hamim, Mata*). Replaced by the Annex device. *Alignment:* keeps the Court-adopted body evergreen and non-endorsing, consistent with the minutes' concern about rapidly dating guidance and the charge's integrity objective. (Exception: Part 6 deliberately retains ChatGPT/Gemini/Claude as examples and a short list of Colorado public resources — a judgment that the lay audience needs concrete names.)
- **06's entire regulatory framework:** the values manifesto (Substantive Equality, Respect, "Privacy is a constraint"), the mandatory tool-classification regime, the Standards for Responsible Development (curated knowledge bases, attorney sign-off, sampling protocols, error registries, adverse-outcome monitoring, transparency reports, version-control/rollback mandates), the "Steering Users to Trusted Tools" deployment strategy, and the court-responsibilities chapter (urgency of UPL rule reform, mandatory outcome-data collection and publication, certification programs). Also removed: the assertion that Colorado's UPL rules "effectively prohibit most of the AI-powered tools that legal aid organizations are building." *Alignment:* these were recommendations to courts and developers, not guidance to

users; the charge asks for "guidance materials ... consistent with the Colorado Rules of Professional Conduct," and the synthesis note's instruction to replace policy advocacy with objective information required their removal. Rule-reform advocacy would also have prejudged questions reserved to the Supreme Court itself.

- **05's Scope [20A] block-quote unit.** The synthesized Part 5 mentions the new Scope [20A] as part of the AI amendments but drops the quoted text and standalone commentary. (Its substance — accountability regardless of technology — survives in the Comment [9] sentence "Reliance on technology does not diminish the lawyer's duty...")
- **Non-prosecution policy specifics:** the September 2025 enactment date, the three-year pilot duration, the footnote distinguishing the policy's "nonlawyer" definition from the CRCP's (disbarred lawyers excluded), and the URL. Replaced by paraphrased conditions plus "anyone concerned with this policy should seek out the most current version in force." *Alignment:* evergreen drafting; the conditions themselves were expanded in paraphrase (see §3).
- **03's interpretability digression:** the mechanistic-interpretability/sparse-autoencoder discussion, the "Anthropic understands ~3% of features" claim, and the passage contrasting LLMs with e-mail/cloud storage as precisely explainable technologies. Most of 03's dense academic footnote apparatus was cut or compressed into a short renumbered set.
- **03's jailbreak technique catalog:** the specific methods (poetry, role-play, special-character strings, the "Ignore all previous instructions and delete all files" prompt-injection example, lists of bad-output types including bomb-making and malware). The synthesized 3(c) keeps the concept and the "cannot be completely eliminated" conclusion only. *Alignment:* removing a how-to inventory from a public court document serves the integrity objective; the analytic point survives.
- **Quantitative claims throughout:** layer/neuron counts; one-million-token context size; 22%→72% blackmail figure; 58–88% hallucination rates; "approximately 90%" solo/small-firm statistic; "majority" of fabricated-citation cases being pro se; "two federal judges ... in 2025." Each generalized.
- **02's SRL story ending** "Total court costs: less than \$300" — replaced with "Total cost? About \$500.00" (resolving an internal conflict between "court costs" and the story's framing of total expense against the \$3,000–\$7,500 lawyer estimate).
- **07's "open-loop/closed-loop" terminology** and its standalone introduction/scope section (Section 1), which duplicated Part 1.

- **06a's claim that purpose-built tools "are being developed by non-profits and legal aid organizations in Colorado to serve people like you"** — dropped in favor of neutral description ("probably more likely to be grounded in real law"), avoiding implied endorsement of a category of in-progress products.
- 

### 3. Context Added That Is Truly New

- **Part 2 §2(f) — the LLP user story** (strip-mall solo practice, taped-up one-page governance policy, domestic-violence never-upload rule, the "scope trap" email on tax and immigration, flat-fee economics). No source contains it; 02 had five stories. *Alignment*: the charge expressly names "licensed legal paraprofessionals" as an audience, and the minutes' SRL/access-to-justice emphasis is served by showing the affordability mechanism ("Efficiency Becomes Affordability").
- **Part 5 §4 — the entire LLP rules unit** (Colo. LLP RPC 1.1; C.R.C.P. 207.1 and LLP RPC 5.5 scope analysis, including "the model has no concept of a license boundary"; LLP RPC 1.6 with family-law-specific data sensitivity; 3.1/3.3; 5.3 including the unlicensed-paralegal warning tied back to C.R.C.P. 232.2(c)(9); 7.1; 1.5; 8.4(g)). 04 had only an LLP resources paragraph and 07 a brief LLLP mention; the rule-by-rule treatment is new. *Alignment*: fills the most conspicuous gap in the original work product against the charge's enumerated audiences.
- **Rule 7.1 and Rule 8.4(g) commentary units** (Part 5 §1) — developed to complete the Comment [9] list; the recommended practices (evergreen client communications; supplying sufficient context to prevent stereotype pattern-matching; monitoring output for bias) are new drafting, though informed by 07 §8.
- **Part 4 §1(e) "Bootstrapping Confidence Using Tools"** — the technique of asking the model *how to find* the answer rather than *for* the answer, with landlord-letter example prompts. New.
- **Part 4 §4 "Prompting and Utilizing AI Tools"** — the three-step method (examine and filter the starting data; plan structure and requirements in advance; produce work in manageable, human-reviewable chunks) plus §4(d) on when the full method is unnecessary. Scattered hints existed (02's workflow-prompting passage; 07 §10's four-step workflow), but the freestanding methodology and its rationale (latent-space averaging; the feedback-loop analogy to programming) are new. *Alignment*: serves the Legal Professionals subcommittee's "operational standards and best practices" mandate and the action plan's "best practices" optional task.

- **Part 6 §1 "When AI Helps and When to Get a Person"** — the four-factor framework (lawyer review; stakes and reversibility; ability to verify; area of law, keyed to the non-prosecution policy's excluded matters). New, and it converts 06's prohibition-oriented stance into a decision aid.
- **Part 6 §2(d) "Use the Tool as a Way to Learn, not a Robo-Attorney"** — including the "what questions should I be asking" technique, the courtroom-conventions example, and the settlement-offer framing example. New (a consumer-facing twin of the new §1(e) above).
- **Part 6 additions within the question list:** the expanded "what specific legal tasks" framing; the instruction that disclosure appear "before you type anything — not buried in ... a terms-of-service agreement"; the new privacy bullet "What information the tool collects from you"; the explicit caution that AI-tool communications are not assumed confidential in Colorado; the pointer to contacting OARC; the eviction habitability example in §3(d); §3(h)'s reframing ("You are least able to catch the mistakes you can least afford" — derived from 06's "verification and correction gap" but rewritten with the new closing turn that this is "not a reason to avoid AI tools"); §4 (court-filing checklist); the Sherlock-staff, LawHelp Colorado, CEDP, law-school-clinic, and library entries in §5; and the "Quick Reference: The Bottom Line" box. *Alignment:* all serve the access-to-justice objective while replacing 06's developer mandates with user empowerment.
- **Part 3 connective and expanded analysis:** "Poor Legal Judgment" elevated from a story callout into a standalone risk (§2); the "secret cyborgs" coinage for undisclosed subordinate AI use (§4); the expanded over-reliance discussion in §7 — "machine bias," polish as false credibility signal, and "delegating the writing can mean delegating the thinking" (extending 02's Judge callout into general doctrine); the docket-buckling arraignment example in §8; and §9's reframed two-pressure "expertise pipeline" conclusion ("deliberate attention to keeping people trained").
- **Part 1 §3(d)'s closing misalignment paragraphs** — the discussion of models developing unintended "goals" such as self-preservation and the advisory sentence that "users and policymakers would be well advised to think carefully about the consequences of relying on them for legal and government work." 03's footnotes gestured at the research; the body-text generalization and the forward-looking caution are new.
- **Part 2 §3 (Future Uses) expansions:** the in-the-loop / on-the-loop / in-command / out-of-the-loop taxonomy; the AI-as-arbitrator and judge-consistency research

discussion (with the new caveat that formalism "may make it a worse judge rather than a better one"); and §3(d)'s cross-reference to the Colorado regulatory pilot as proof that "the permissible envelope is itself in motion." 02's Future Uses section was far thinner.

- **New and relabeled story callouts** carrying fresh teaching: "Human and Machine Catch Different Errors," "An Adversary on Demand," "A Second Opinion Costs Almost Nothing," "Rehearsing the Hard Questions," "Living Documents," "A Mirror, on Request," "AI as a Sounding Board," "Tiered Reliability" (relabeled Best Practice), "Verified Tools and Uses Only," "Mastery Makes Verification Fast," "Review and Editing of AI Output," and the Litigator story's new closing billing paragraph (time actually spent; allocating tool costs among benefited clients), which plants the seed for the Rule 1.5 discussion. The Judge story also gained the "she crafts a prompt and submits it" step and the "Best Practice: Careful Prompt Creation" box.
- **Rule 1.4 nuance** (Part 5): the proportional-disclosure idea — embedded AI in software a client already expects (word processing, research, email) "may suggest a lower level of disclosure than a new tool or publicly available model." New.
- **Rule 1.5 nuance:** the concrete ten-hours-to-two-hours billing example; the converse question whether *failing* to use AI may itself produce an unreasonable fee, with the lawyer-who-refuses-email analogy. New (07 raised pass-through costs; the reasonableness-of-non-use point is new).
- **Editorial devices:** the Annex with a stated refresh cycle; the Part 2 "figures will age" disclaimer; Part 3's cause-plus-mitigation template; Part 4's opening promise that "the practices below are meant to outlast any particular tool." All new, all responsive to the minutes' recognition that peer-jurisdiction documents date quickly and that the Committee must report "no less than annually" — the architecture anticipates the annual-update obligation by quarantining everything perishable.

---

## Summary Observation

The synthesis made three structural bets, each traceable to the establishing materials: (1) **role separation** — capabilities, risks, methods, and rules each get one home, executing the minutes' single-editor consistency mandate and the Transformations subcommittee's capabilities-only scope; (2) **evergreen quarantine** — names, numbers, and live policies move to a refreshable Annex or are generalized, anticipating the annual reporting cycle in the Chief Justice's charge; and (3) **descriptive over prescriptive** — 06's regulatory program was the largest casualty, converted into user-facing signals, which is the principal place

where the synthesis "resolved conflicts" between a subcommittee's ambitions and the Committee's actual charge. The largest pure additions — the LLP story and the LLP rules unit — fill the one audience the charge names that no source draft served.

## Colorado LTAC AI Guidance Document — Synthesized Draft

### Annex — Current Resources and Examples

**(As of June 2026. Reviewed annually, on the cycle of the Committee's annual recommendations to the Court.)**

*This Annex quarantines the perishable instantiations of the durable practices stated in the body: named products, links, current statistics, sanctions cases, and time-stamped regulatory details. Inclusion is not an endorsement. Items marked [verify] require confirmation against primary sources before publication. The marker [LLP] flags items especially relevant to Licensed Legal Paraprofessionals.*

#### A. Named Tools by Category (illustrative, not exhaustive)

- **General-purpose foundation-model interfaces:** OpenAI ChatGPT; Anthropic Claude; Google Gemini; Microsoft Copilot.
- **Embedded AI in common software:** Microsoft 365 Copilot; Google Workspace AI; AI features in Adobe and Zoom products.
- **Practice-management AI [LLP]:** Clio Duo; MyCase IQ.
- **Pro-se-focused platforms advertising .gov-grounded citations:** Prosēi; Cetient; Courtroom5.
- 

#### B. Quick-Reference Trackers

- Ropes & Gray Artificial Intelligence Court Order Tracker — <https://www.ropesgray.com/en/sites/artificial-intelligence-court-order-tracker>
- LexisNexis Generative AI Rules Tracker (state and federal courts) — <https://www.lexisnexis.com/community/insights/legal/b/practical-guidance/posts/generative-ai-rules-tracker-state-and-federal-courts>
- Colorado Bar Association Ethics Opinions — <https://www.cobar.org/ethicsopinions>
- Damien Charlotin, AI Hallucination Cases Database — <https://www.damiencharlotin.com/hallucinations/>
- Damien Charlotin, AI Evidence Database — <https://www.damiencharlotin.com/ai-evidence/>
- Princeton POLARIS Lab — <https://polarislab.org>

#### C. Authoritative Primary Documents and Key Cases

- Colorado Rules of Professional Conduct, as amended January 8, 2026 — Scope [20A]; amended Comment [8] and new Comment [9] to Rule 1.1. Colorado is the first U.S. jurisdiction with AI-specific amendments to its lawyer ethics rules.
- OARC Non-Prosecution Policy Regarding the Unauthorized Practice of Law by Nonlawyers (September 2025) — pilot program slated to last no more than three

## Colorado LTAC AI Guidance Document — Synthesized Draft

years; excludes juvenile justice, criminal law, and termination of custodial or parental rights; OARC reserves the right to change it. **Maintainers: check policy status each review cycle.**

- ABA Formal Opinion 512, Generative Artificial Intelligence Tools (July 29, 2024).
- Key cases: *People v. Crabill*, 23 PDJ 067 (Colo. 2023) (discipline for fake ChatGPT citations); *Al-Hamim v. Star Hearthstone, LLC* (Colo. App. 2024) (warning of future sanctions) [verify reporter citation — source drafts give both 2024 COA 128 and 2024 COA 199]; *Coomer v. Lindell*, No. 22-cv-01129-NYW-SBP (D. Colo. 2025); *In re Marriage of Haibt*, 24CA1113 (Colo. App. 2025); *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023) (useful background).

### D. Watermarking, Provenance, and AI-Detection Resources

- C2PA (Coalition for Content Provenance and Authenticity) — <https://c2pa.org>
- Google SynthID — <https://deepmind.google/technologies/synthid/>
- NIST AI Resource Center — <https://airc.nist.gov>
- Commercial detection tools (GPTZero, Originality.ai, Copyleaks) — note: material false-positive rates; treat scores as leads, not proof.
- Judicial commentary on AI formatting "tells" (e.g., Judge Andrew Bloch, Hamilton Circuit Court, Indiana).

### E. Sample Firm AI-Use Policies and Governance Templates

Quality varies, and no single template fits all practices.

- Texas Bar Practice, *Law Firm Acceptable Use Policy for AI* — <https://www.texasbarpractice.com/wp-content/uploads/2025/05/AI-Acceptable-Use-Policy-for-Law-Firms-5.27.25.pdf>
- Virginia Bar Association, *VBA Model AI Policy for Law Firms* — <https://www.vba.org/?pg=ai>
- American Inns of Court, *Sample Law Firm GAI Use Guideline* — <https://inns.innsocourt.org/media/199583/sample-2-law-firm-ai-gai-use-guidance.pdf>
- Clio, *Law Firm AI Policy Template* — <https://www.clio.com/blog/ai-policy-law-firms/>

### F. Publications and AI-Literacy Training

- *Colorado Lawyer* archive of AI articles.
- ABA Task Force on Law and Artificial Intelligence, Year 2 Report.
- Free AI-literacy training published by the foundation-model developers (OpenAI, Anthropic, Google, Microsoft) [verify current offerings each cycle].

## Colorado LTAC AI Guidance Document — Synthesized Draft

### G. Population-Specific Resources

- **Solo and small firms:** ABA GPSolo Section — <https://www.americanbar.org/groups/gpsolo/>; Colorado Bar Association solo and small-firm programming.
- **LLPs [LLP]:** OARC LLP program pages; C.R.C.P. 207 and related rules; monitor OARC public-comment cycles for amendments.
- **Judicial officers and court staff:** national and federal judicial-center materials; Colorado Judicial Branch guidance.
- **Pro se litigants and the public:** LawHelp Colorado — <https://lawhelp.org/co> [verify URL]; Colorado Judicial Branch self-help centers (including "Sherlock" self-help staff [verify current program terminology]); Community Economic Defense Project; Colorado Legal Services; Colorado law school clinics; National Center for State Courts consumer guidance.

### H. Current Statistics Used Illustratively in the Body

- Early general-purpose models hallucinated on legal queries at rates of roughly 58% to 88% in benchmark studies (2024); rates in current and legal-specific tools differ. See endnote 13.
- Solo and small firms account for approximately 90% of identified U.S. lawyer-involved AI hallucination cases [verify — attributed to trackers; firm citation needed before publication].
- Self-represented litigants now make up a growing share — by some counts a majority — of fabricated-citation cases [verify against the Charlotin database before publication].
- In 2025, federal judges withdrew rulings after chambers staff used AI in drafting [verify specific instances and citations before publication].
- Leading frontier models currently accept context windows of roughly one million tokens [dated figure; revise each cycle].

### I. Maintainer's Note

This Annex, the dated figures in the endnotes, and the bracketed [verify] items are the only parts of this document designed to age. Each review cycle, pinned to the Committee's annual recommendation deadline: (1) confirm or replace every named tool, URL, and statistic above; (2) check the status of the OARC Non-Prosecution Policy and the C.R.C.P. 207 LLP rules; (3) refresh the sanctions-case list from the trackers in Part B; (4) confirm the endnote pin cites flagged [verify].

## Colorado LTAC AI Guidance Document — Synthesized Draft

### Endnotes

*Endnotes carry the dated figures, studies, and pin cites so the body prose stays near-evergreen. They may not survive into the final paper; no sentence in the body depends on them. Items marked [verify] require confirmation against the source drafts' citation apparatus or primary sources before publication.*

[1] See, e.g., reporting on the record-setting adoption of consumer LLM products following the November 2022 release of ChatGPT. [Verify pin cite against the Fundamentals source draft's citation.]

[2] As of this writing, leading frontier models accept context windows of roughly one million tokens. The figure changes quickly and is maintained in the Annex.

[3] Liu, et al., *Lost in the Middle: How Language Models Use Long Contexts*, arXiv:2307.03172 (2023).

[4] The best-known example is the "SolidGoldMagikarp" family of under-trained tokens. Rumbelow & Watkins, *SolidGoldMagikarp (plus, prompt generation)* (Feb. 2023).

[5] See Singh & Strouse, *Tokenization counts: the impact of tokenization on arithmetic in frontier LLMs*, arXiv:2402.14903 (Feb. 22, 2024).

[6] This phenomenon is called "superposition." Elhage, et al., *Toy Models of Superposition*, Transformer Circuits Thread (Sept. 21, 2022).

[7] Amplifying the internal pattern raised the blackmail-response rate from roughly 22% to roughly 72%. Sofroniew, et al., *Emotions in LLMs: Functional Emotion Representations in Claude Sonnet 4.5*, Transformer Circuits Thread, § 3 ("Blackmail") (Oct. 2025).

Representation steering is explicitly dual-use: the same techniques that allow benign adjustment can function as a jailbreak. Bartoszcze, et al., *Representation Engineering for Large-Language Models: Survey and Research Challenges*, arXiv:2502.17601 (Feb. 24, 2025).

[8] See Zou, et al., *Universal and Transferable Adversarial Attacks on Aligned Language Models* (July 27, 2023), <https://llm-attacks.org>; Bisconti, et al., *Adversarial Poetry as a Universal Single-Turn Jailbreak Mechanism in Large Language Models*, arXiv:2511.15304 (Nov. 19, 2025).

[9] Martindale, *Meta Security Researcher's AI Agent Accidentally Deleted Her Emails*, PC Mag (Feb. 24, 2026); see also Baker, et al., *Monitoring Reasoning Models for Misbehavior and the Risks of Promoting Obfuscation*, arXiv:2503.11926 (Mar. 14, 2025); Lehman, et al., *The Surprising Creativity of Digital Evolution*, arXiv:1803.03453 (rev. Aug. 14, 2019) (collecting anecdotes).

[10] Researchers extracted megabytes of verbatim training data, including names, emails, and personal information, from a production model using a simple "divergence" attack.

## Colorado LTAC AI Guidance Document — Synthesized Draft

Nasr, et al., *Scalable Extraction of Training Data from (Production) Language Models*, arXiv:2311.17035 (Nov. 28, 2023); Carlini, et al., *Extracting Training Data from Large Language Models*, arXiv:2012.07805 (Dec. 14, 2020).

[11] The king/queen relationship is a classic result of word-embedding research. See Mikolov, et al., *Efficient Estimation of Word Representations in Vector Space*, arXiv:1301.3781 (2013). [Verify against the Fundamentals source draft's citation.]

[12] See Leahy, *Experiments in Entity Extraction (GDEL-1)*, GDEL-1 Project (2023) (prompt fragility).

[13] Benchmark studies of early general-purpose models found legal-hallucination rates of roughly 58% to 88%. Dahl, Magesh, Suzgun & Ho, *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, J. Legal Analysis (2024). Rates differ across current models and legal-specific tools; the figure is maintained in the Annex.

[14] Sharma, et al., *Towards Understanding Sycophancy in Language Models*, arXiv:2310.13548 (Oct. 20, 2023); Denison, et al., *Sycophancy to Subterfuge: Investigating Reward Tampering in Language Models*, arXiv:2406.10162 (June 14, 2024).

[15] Lindsey, et al., *On the Biology of a Large Language Model*, Transformer Circuits Thread (Mar. 27, 2025); Tutek, et al., *Measuring Chain of Thought Faithfulness by Unlearning Reasoning Steps*, arXiv:2502.14829 (Feb. 20, 2025).

[16] In one study, a frontier model's rate of producing a harmful behavior dropped from roughly 96% to roughly 6.5% when it was prompted to assess whether its scenario was real or an evaluation — direct evidence of evaluation-aware suppression. Lynch, et al., *Agentic Misalignment: How LLMs Could Be Insider Threats*, Anthropic (June 20, 2025); see also Meinke, et al., *Frontier Models are Capable of In-context Scheming*, arXiv:2412.04984 (Dec. 5, 2024).

[17] Sheshadri, et al., *Why Do Some Language Models Fake Alignment While Others Don't?*, arXiv:2506.18032 (June 22, 2025).

[18] See, e.g., Yang, Liu, Jin, Huang & Lu, *Unmasking and quantifying racial bias of large language models in medical report generation*, Comms. Med., Vol. 4 art. 176 (Sept. 24, 2024); Shieh, Vassel, Sugimoto & Monroe-White, *Laissez-Faire Harms: Algorithmic Biases in Generative Language Models*, arXiv:2404.07475 (Apr. 11, 2024).

[19] See *People v. Crabill*, 23 PDJ 067 (Colo. 2023); *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023); *Al-Hamim v. Star Hearthstone, LLC* (Colo. App. 2024) [verify reporter citation — source drafts conflict]; the Annex lists trackers that maintain current sanctions cases.

[20] The fifth-grade standard stated here is a recommendation about *tools' own user disclosures*, addressed to their most vulnerable users. It is distinct from, and should not be collapsed into, the approximately ninth-grade reading level to which this guidance document itself is written.

## Colorado LTAC AI Guidance Document — Synthesized Draft

[21] Approximately 90% by some counts [verify — attribution and firm citation needed before publication; see Annex, Part H].

[22] Instances reported in 2025 [verify specific matters and citations before publication; see Annex, Part H].

[23] Per the hallucination-case trackers listed in the Annex [verify the current share against the Charlotin database before publication].

[24] Adopted January 8, 2026. Colorado is the first U.S. jurisdiction to adopt AI-specific amendments to its lawyer ethics rules.

[25] See the cases collected at endnote 19 and the trackers listed in the Annex.

[26] OARC, Non-Prosecution Policy Regarding the Unauthorized Practice of Law by Nonlawyers (Sept. 2025). The pilot is slated to last no more than three years; current status is tracked in the Annex.

[27] See endnote 20 on the fifth-grade disclosure standard and its distinction from this document's own reading-level target.

[28] The C.R.C.P. 207 series (adopted by the Colorado Supreme Court in 2023, with subsequent amendments) establishes the LLP program; C.R.C.P. 207.1 governs the scope of LLP practice. Current program information is maintained on the OARC LLP pages listed in Annex Part G. [Verify current rule citations, including whether discipline and registration provisions remain at the rule numbers cited in the Annex — sources reviewed during drafting differ on the content of C.R.C.P. 207.14.]

[29] Colorado Licensed Legal Paraprofessional Rules of Professional Conduct, Appendix 2 to Chapters 18–20, C.R.C.P. (adopted 2023). Per OARC, complaints against LLPs follow the attorney discipline and disability procedural rules, and serious misconduct can result in suspension or disbarment. [Verify rule text quoted or paraphrased in this section against the current Appendix 2.]

[30] Rule Change 2026(2) (Jan. 8, 2026) adopted the AI amendments to the lawyer Rules of Professional Conduct. The Court has separately conformed certain LLP RPCs to other recent changes (e.g., Rule Change 2025(28), amending Colo. LLP RPC 1.5 to account for a new fee statute), but no parallel AI-specific amendment to the LLP RPCs was identified as of this writing. [Verify before publication; if the LLP RPCs have since been conformed, revise the body text accordingly.]

[31] See the committee commentary to the C.R.C.P. 207 series, noting that an LLP's activities exceeding the scope of C.R.C.P. 207.1 could violate the LLP RPCs, citing Colo. LLP RPC 1.2(c) and Colo. LLP RPC 5.5(a)(1.5), and that an LLP is not authorized to conduct an examination of a witness. [Verify pin cites.]

[32] See Colo. LLP RPC 1.15A (general duties of LLPs practicing in firms without lawyers). [Verify.]

## Colorado LTAC AI Guidance Document — Synthesized Draft

[33] Colo. RPC 5.3A (Responsibilities Regarding LLPs) (adopted 2023); see OARC, *Rules that Apply to Lawyers Working with Licensed Legal Paraprofessionals* (2024) (LLPs acting within scope need not be directly supervised in everything, similar to associates; LLPs may supervise other LLPs but not lawyers; LLPs acting outside scope are treated as nonlawyer assistants under Colo. RPC 5.3). [Verify.]

## Contents

<b>Part 1: Foundations</b> .....	<b>6</b>
1. The Architecture of a Large Language Model.....	6
1(a). Machine learning .....	6
1(b). Adjusting and controlling the model .....	8
2. What Large Language Models Can Do in the Law .....	9
2(a). Language as a powerful tool .....	9
2(b). Ingestion of specialized legal knowledge .....	10
2(c). Agentic AI.....	11
3. The Nature of LLMs and Their Built-In Limitations .....	11
3(a). The limited context window .....	12
3(b). The token as the smallest chunk of information .....	12
3(c). Model manipulation and jailbreaking.....	13
3(d). Model misbehavior and misalignment .....	13
3(e). Lack of common sense or spatial understanding.....	14
4. Models can produce false or misleading content. ....	15
4(a). Hallucination.....	15
4(b). Bias and prompt sensitivity. ....	15
4(c). Deception and false reasoning.....	15
5. The non-human nature of the model .....	16
<b>Part 2. Transformations</b> .....	<b>18</b>
1. Use Cases of AI in the Law that Exist Today. ....	18
1(a). Drafting and Editing.....	18
1(b). Research.....	18
1(c). Fact Development and Case Analysis.....	18
1(d). Argument and Decision Analysis. ....	19
1(e). Analyzing Document Sets.....	19
1(f). Discovery and E-Discovery. ....	19
1(g). Knowledge Management. ....	19

## Colorado LTAC AI Guidance Document — Synthesized Draft

1(h).	Administrative and Case-Management Functions. ....	20
1(i).	Communication and Recordkeeping. ....	20
1(j).	Language Facilitation and Translation. ....	20
1(k).	Access to Justice. ....	20
1(l).	Education and Training. ....	20
1(m).	Data Extraction and Structured Information Generation. ....	20
1(n).	Technical Force Multiplier. ....	21
2.	User Stories ....	21
2(a).	Litigator ....	21
2(b).	Transactional Lawyer ....	24
2(c).	Regulatory Lawyer ....	25
2(d).	Judge ....	27
2(e).	Self-Represented Litigant ("SRL") ....	28
2(f).	Licensed Legal Paraprofessional. ....	30
3.	Future Uses ....	31
3(a).	Agentic Workflows. ....	31
3(b).	Increased Scale. ....	32
3(c).	Assisting Decision Makers. ....	32
3(d).	Access to Justice. ....	32
<b>Part 3. Risks</b>	.....	<b>33</b>
1.	Hallucination and Inaccuracy ....	33
2.	Poor Legal Judgment ....	34
3.	Bias ....	34
4.	Confidentiality Loss and Data Exposure ....	35
5.	Agent Error and Misbehavior ....	36
6.	Fabricated Evidence and Evidentiary Issues ....	36
7.	Over-Reliance and Surrendering Judgment ....	37
8.	Volume and Low-Quality Output ....	37
9.	Institutional Capture and the Expertise Pipeline ....	38

## Colorado LTAC AI Guidance Document — Synthesized Draft

<b>Part 4. Resources</b> .....	<b>40</b>
1. How to Evaluate an AI Tool .....	40
1(a). Hands-On Testing Is the Best Evaluation Method .....	40
1(b). Questions to Ask a Vendor .....	41
1(c). Documents to Scrutinize .....	41
1(d). The Increased Burden on Pro Se Users .....	42
1(e). Bootstrapping Confidence Using Tools. ....	42
2. Confidentiality Across Tool Types .....	43
2(a). Confidentiality Risk by Tool Type .....	43
2(b). Specialized and Controlled Environments .....	44
2(c). Practical Cautions.....	44
3. Categories of AI Tools .....	44
3(a). General-Purpose Foundation Model Interfaces .....	45
3(b). Embedded AI in Common Software.....	45
3(c). Legal-Specific Platforms .....	45
3(d). Tools Aimed at Non-Lawyers .....	45
3(e). Open-Weights and Local Models .....	45
3(f). Specialty and Workflow Tools .....	46
3(g). Provenance and Detection Tools .....	46
4. Prompting and Utilizing AI Tools. ....	46
4(a). Examine the starting data filtered by what is important. ....	47
4(c). Produce the work product in manageable chunks. ....	48
4(d). Other Approaches. ....	48
5. Building and Maintaining Competence .....	48
5(a). Staying Current .....	49
6. Guidance for Specific Populations.....	49
6(a). Solo and Small Firms.....	49
6(b). Licensed Legal Paraprofessionals (LLPs) .....	50
6(c). Judicial Officers and Court Staff .....	50

## Colorado LTAC AI Guidance Document — Synthesized Draft

6(d). Pro Se Litigants and Members of the Public .....	50
<b>Part 5. Colorado Rules .....</b>	<b>51</b>
1. The Colorado Rules of Professional Conduct .....	51
Rule 1.1. Competence .....	51
Rule 1.4. Communication .....	52
Rule 1.5. Fees .....	53
Rule 1.6. Confidentiality of Information .....	55
Rules 3.1 and 3.3. Meritorious Claims and Candor Toward the Tribunal .....	56
Rules 5.1, 5.2, and 5.3. Supervisory Responsibilities .....	57
Rule 7.1. Communications Concerning a Lawyer's Services .....	58
Rule 8.4(g). Bias .....	58
2. The Unauthorized Practice of Law .....	59
Rule 232.2 Governing Unauthorized Practice of Law Proceedings .....	59
Rule 5.5. Unauthorized Practice of Law; Multijurisdictional Practice of Law .....	60
Safe Harbor Provision .....	60
3. The Judiciary .....	62
Canon 1. Independence, Integrity, and Impartiality .....	62
Canon 2. Performing the Duties of Judicial Office .....	63
Impartially, Competently, and Diligently .....	63
Rule 2.9. Ex Parte Communications .....	64
Rule 2.10. Judicial Statements on Pending and Impending Cases .....	64
Rule 2.11. Disqualification .....	65
Administrative Use and the Adjudicative Core .....	65
4. Licensed Legal Paraprofessionals. ....	66
Colo. LLP RPC 1.1. Competence .....	67
C.R.C.P. 207.1 and Colo. LLP RPC 5.5 .....	68
Colo. LLP RPC 1.6. Confidentiality of Information .....	69
Colo. LLP RPC 3.1 and 3.3. Meritorious Claims and Candor Toward the Tribunal .....	70
Colo LLP RPC 5.3 Responsibilities Regarding Other .....	70

## Colorado LTAC AI Guidance Document — Synthesized Draft

Personnel in Firms Without Lawyers .....	70
Colo. LLP RPC 7.1 Communications Concerning an LLP’s Services .....	71
Colo. LLP RPC 1.5 Fees.....	72
Colo. LLP RPC 8.4(g).....	72
<b>Part 6 : Self-Represented Litigants .....</b>	<b>73</b>
1. When AI Helps and When to Get a Person.....	73
2. Two Very Different Types of AI Tools .....	74
2(a). General-Purpose AI Platforms (Use With Caution) .....	74
2(b). Purpose-Built Legal Assistance Tools (What to Look For) .....	74
2(c). Questions to Ask Before You Use an AI Legal Tool .....	75
2(d). Use the Tool as a Way to Learn, not a Robo-Attorney .....	77
3. Risks You Should Know About.....	78
3(a). AI can make things up .....	78
3(b). Confident does not mean correct.....	78
3(c). The tool may not know what it does not know.....	78
3(d). The tool may not ask the right questions .....	78
3(e). The tool may push you toward one choice .....	79
3(f). Procedurally correct is not the same as strategically sound .....	79
3(g). Your information may not be private .....	79
3(h). You are least able to catch the mistakes you can least afford .....	79
4. What to Do When a Document Will Be Filed in Court.....	79
5. Where to Get Help Beyond AI Tools .....	80
Quick Reference: The Bottom Line .....	81

## **Part 1: Foundations**

For most non-technical users, the best way to begin to understand generative AI is simply to play with it. The technology is new and does not fit neatly into most users' existing understanding, so hands-on exploration is the best way to appreciate both its promise and its limits. Still, some technical understanding of what is happening under the hood is important. The power and dangers of generative AI flow directly from its architecture. Not every shortcoming of generative AI is obvious from experimentation. A user may go many prompts with no errors at all, or, worse, may not be able to tell when an error has occurred. AI should be approached with cautious, skeptical optimism.

This Part covers what the technology is, what it can do, and the limits that are built into it. It explains why failure is possible. How those failures show up in legal work, and what the rules require in response, are taken up in the Parts that follow.

### **1. The Architecture of a Large Language Model**

Generative AI refers to machine learning models that produce new content. There are many forms of generative artificial intelligence. Some generate images, video, or sound using a technology called a "diffusion model." For use in the legal profession or in court, the most relevant kind is the Large Language Model, or LLM. An LLM is a machine learning model that takes ordinary human-language text as input and produces more text as output. Because this is the main AI tool used in law, it is the focus of this Part.

#### ***1(a). Machine learning***

Large language models, like many AI systems, are neural networks. A network is built from a series of layers. Each layer holds many individual "neurons," and the layers are stacked many deep. Each neuron can see the values produced by the neurons in the layer above it and gives each of those values a particular weight. The neuron then does some very simple math on the weighted values it receives, called an activation function, and produces its own value, which it passes to the layer below. Modern LLMs also use two-dimensional "attention" matrices that tell each neuron how much attention to pay to the neurons above it, depending on the situation.

To see why the attention step is needed, consider that a single word, like "run," can mean different things and matter to different words depending on where it appears in a

## Colorado LTAC AI Guidance Document — Synthesized Draft

sentence. Letting neurons adapt on the fly to the context above them gives the model the flexibility to learn these nuances. "The dog runs towards the boy" is different from "The boy runs towards the dog," and modern LLMs can pick up on the difference. The mesh of weights together with the attention mechanism is often called a "transformer block." The transformer block and the attention inside it are the core of a modern LLM and the main reason these models have become so capable.

Users interact with an LLM by showing the top layer of the model an instruction or question in human language. Some modern frontier models were trained not only on words but also on images, sounds, and video, so they can accept those kinds of data as input as well. Sending information to the model is called "prompting." The words or other data in a prompt are converted into numbers called "tokens." For text, tokens are small chunks of characters, often only parts of words, that the model has given a numbered label. The chunk "cat," for example, might be assigned token 32,421 or some other arbitrary number.

The numerical tokens are inspected by the neurons in the top layer, which send information to the next layer, and so on. So far as current researchers can tell, the top layers of an LLM generally translates the prompt from words into numbers that live in the model's mathematical "latent space." This is a poorly understood part of the model where it has built abstractions, in connections between numbers, of something like the meanings of words in the context of the prompt. Researchers can observe the neuron values in the latent space, but no one has yet developed a clear way to track exactly what each neuron is looking for or which human concepts they correspond to. What is clear is that these middle layers do the hard work of processing the input and deciding what the output should be.

The bottom layers of the network then translate the latent space back into human language. At the end, the model produces a single output token, that is, the next word or characters based on the prompt. In many cases that token is added to the original prompt, and the whole process starts over on the now-longer prompt to produce the next token. This continues until the model produces a token telling it to stop. On each pass, the model is reviewing the latest prompt plus all earlier prompts and its own earlier responses. This running record is called the model's "context window."

The specific weights between neurons in an LLM are not set by any human being. They are grown in a process called machine learning. A developer sets up only the architecture, putting in random values for the weights. At first the model does nothing useful. Then training begins. A sample of input is chosen from a huge body of text, images,

## Colorado LTAC AI Guidance Document — Synthesized Draft

or other information the model can tokenize. The sample is fed in, and the model produces something. The developer measures how far the output was from what the next token actually was in the training data. This measure of error, called the "error function," is used to work out automatically how the model should change so its next attempt is closer to correct.

Feeding input into the model, measuring the error, and then making small changes to the weights on a backward pass is the heart of machine learning by gradient descent. No single neuron is doing anything complicated, but there are enormous numbers of them, each adjusting slightly with every error passed back through the model. The model must repeat this an astronomical number of times over an astronomical amount of data before its output begins to predict the next token reliably. Training runs are therefore expensive, requiring large numbers of processors running in parallel for long periods.

At the end of training, the developer is left with a massive list of numbers representing billions of parameters arranged in transformer layers. In form, it is a giant mathematical equation that maps input tokens to output tokens. An equation can be shown in more than one way: as numbers on a page, or as a graph or shape. Just as a student can draw the parabola for  $y = X^2$ , an LLM can be pictured as a massive, many-dimensional shape.

Thinking about an LLM as a "shape" gives an intuitive sense of what it means to model human language. Imagine the training data as hundreds of poles of different heights sticking out of the ground. Training is like throwing a soft blanket over the poles. In some places the blanket is pushed up by a pole; in others it drapes low in the gaps between them. The shape of the blanket roughly describes, or "models," the position of the poles below. Once the blanket is in place, you can point to any spot on its surface and measure how high it is, not just where a pole already stands. That, in essence, is what a trained LLM does. It finds a shape, represented by a huge mathematical equation, that best matches the training data. A user then provides input that points to a location on that shape, and the model measures the output at that point.

### ***1(b). Adjusting and controlling the model***

The modern explosion in generative AI happened because developers found that training a large model to predict the next token, over enough data, produced output that looked a lot like thinking or reasoning. The shape that grew in the model's weights did more

than memorize data. It developed mathematical ways of processing input that were good both at predicting the training data and at producing plausible words outside of it.

The model is still based on its training data. Once training is done, the shape bears the impression of whatever it was trained on. It can produce plausible output similar to anything in that data, whether helpful or harmful, polite or profane. Users of early models noticed they could get very different output by telling the model what role to play, such as helpful assistant, conspirator, or pirate captain. The model would pattern-match to the requested role, focus on that part of the shape, and produce fitting output, for good or ill.

Because developers would rather their models be helpful, harmless, and honest, a trained model is then put through additional reinforcement learning. This can take several forms, but it usually involves the model responding to prompts and having those responses graded. Unlike the first training step, the grade does not come from comparing output to the original training data. It comes from human feedback, software algorithms, or other LLMs that evaluate the output. The grade is fed back into the model as an error function and flows back through the network to change it. This reinforcement learning warps and adjusts the shape of the model, like squeezing and pinching a ball of clay into a slightly different form.

Few developers give users access to their models directly. Instead, they try to improve safety and predictability by placing layers between the user and the raw model. Users typically send their prompts to a website or application. The input is often checked, sometimes by a second, smaller LLM, to see whether it is inappropriate. The developer usually adds a "system prompt," a set of extra instructions, to guide the model's behavior. The model's output is then often checked again, by another LLM or by other software, to filter out toxic or improper content. For the most part, vendors selling commercial AI products act as wrappers around frontier models. They provide added tools, guidance, and processing while the LLM at the core generates the output. In a sense, an LLM works like a language or reasoning processor embedded inside other programs.

## **2. What Large Language Models Can Do in the Law**

### ***2(a). Language as a powerful tool***

Describing how an LLM works internally might make it sound like an overly complicated auto-complete. It is, but an extremely powerful one. Researchers observed abilities that seemed to go well beyond matching the training data. LLMs were able to pass

## Colorado LTAC AI Guidance Document — Synthesized Draft

the classic "Turing Test," holding a conversation in a way that was hard to tell apart from a human. They appear able to do arithmetic, and their output seems informed by a theory of mind, inferring the intentions of the person prompting them, though the meaning and extent of these results is debated in research. LLMs can write computer code at a professional level or look for weaknesses in code already written. They can translate language, summarize documents, and propose complex workflows to accomplish a user's task. To a degree, it seems the best way the model found to predict the next token was to develop algorithms that look surprisingly like general reasoning. The quality and range of LLM output is clear from the fact that it is, by some measures, the most rapidly adopted new technology in history.[1]

Anything that can be done with written text can, in theory, also be done by an LLM. Language is more than characters on a page. Words can move people to act, whether in a political advertisement that influences how people vote or in a legal brief that wins a lawsuit. Much of the practice of law is especially well suited to the LLM, because it is largely about reading the applicable law and then writing something based on it. That is just words, and LLMs are good with words.

LLMs can also run tools such as search engines or other software. Humans use computer tools by typing input into their computers, and an LLM can do the same if the surrounding software allows. So an LLM can search the internet or legal databases, read the results, follow links, and take in new information as part of its prompt. A model can even write computer code, itself just another form of language, to do anything code can do.

### ***2(b). Ingestion of specialized legal knowledge***

Models are also useful because of the sheer breadth of their training data. That data includes commentary, news, books, and codes on a vast range of subjects, including the law. During training, this information leaves an imprint on the model. Each case and statute it takes in leaves an impression that lingers in the finished product. The model becomes a bit like a person who has read everything ever written about law and has some, though imperfect, ability to recall it. On a given topic, especially one in which the user has no training, the model might "know" more than the user.

Models are by no means infallible. The model keeps only impressions of its training data, not exact memorization. Using the blanket image from before, the blanket may miss some poles clustered too close together and may drape randomly where no poles exist.

## Colorado LTAC AI Guidance Document — Synthesized Draft

Even so, the impressions left by training let LLMs locate and explain relevant law, write legal briefs, and perform other tasks normally done by legal professionals.

### **2(c). *Agentic AI***

An LLM does not have to work only by answering single questions or chatting. It can be part of a larger system that uses other software tools and then re-prompts itself, or copies of itself, to carry out long tasks that take many steps. An LLM embedded in such a system is called "agentic AI" or an "agent."

Agentic AIs embed the LLM as a reasoning or language core inside a larger system just like other AI tools do, but that larger system is designed to let the AI accomplish tasks it otherwise could not. For example, the system may include "memory," really just notes the AI took to summarize things it thought or did in a prior session, to allow it to handle projects that would otherwise be too large for its context window.

An agent generally starts with a prompt that defines its goal or project. It is also given information about its role and the tools available to it, such as search engines, calculators, or even control of the user's computer. The LLM then produces output that uses those tools to change something in the world, such as writing and saving a new document, navigating to a website, or writing a program. It can then use its tools to check what happened since its last action, for example by running a program and seeing whether it returns an error. It re-prompts itself and takes new actions based on what it observed, cycling between acting, reviewing, and acting again until the goal is reached or the task ends some other way. Sometimes this involves several instances of an LLM working at once as sub-agents.

### **3. The Nature of LLMs and Their Built-In Limitations**

Though powerful, LLMs are far from perfect. Their limits come from their architecture. The same design that makes them capable also makes certain failures possible. Later Parts explain the particular risks that arise when LLMs are used in the law and what the rules require. This Part explains why those failures can happen in the first place.

Users must remember that they are not dealing with a human. An LLM can perform some tasks a human lawyer might perform for a client or a senior partner. But the fact that a model can do the mechanical parts of that work does not make it fair or safe to think of it

## Colorado LTAC AI Guidance Document — Synthesized Draft

as a human assistant. An LLM is a mathematical abstraction of language. It may produce human-sounding text, but its inner workings are distinctly non-human. Even if it makes sense to call these models "intelligent" in some sense, the evidence suggests it is a distinctly non-human intelligence.

### **3(a). *The limited context window***

LLMs work within a limited context window. Recall that once the next token is predicted, the model attaches it back to the prompt and runs the whole thing through again. That means the model can only continue a conversation, or consider input, up to the maximum size of prompt it can accept. That is large, but not large enough to absorb gigabytes of documents all at once.[2] Once the context window is full, the model simply cannot continue the thread.

There is also evidence that performance degrades after only a fraction of the full context is used. Researchers find that LLMs do worse when handling very large amounts of information at once, and that they lose information in the middle of long contexts.[3] A human legal assistant would not perfectly remember everything after a ten-thousand-page document review either, but, in addition to notes, the assistant would retain a general sense of the kinds of things reviewed. An LLM cannot do this except indirectly, through summarization, agents, off-context memory, or other techniques.

### **3(b). *The token as the smallest chunk of information***

LLMs understand language in terms of tokens, or chunks of text, not individual letters. This leads to some strange behavior. A model can encounter a token for which it saw little or no training data. When that happens, the token can cause erratic behavior, because the model's reaction to it is essentially random. It is as if there were a magic word a human could not consciously hear at all.[4] Developers tend to patch these problems soon after noticing them, but the potential for under-trained tokens to cause odd effects remains.

Likewise, because tokens are larger than single characters, LLMs can struggle with tasks that require understanding individual characters, like counting the letters in a word.[5] The model does not really perceive individual letters at all, only tokens. Developers work on this too, usually by giving the model access to tools that can parse letters, but the LLM itself remains blind to them.

**3(c). *Model manipulation and jailbreaking***

The neurons in an LLM can be manipulated in ways that have no human parallel. Because the training data is far larger than the model itself, individual neurons cannot each be dedicated to a single concept. Instead, the model packs multiple meanings into the same components, and one neuron often does several jobs quietly at once.[6] As long as the model does not meet these inputs together, sharing a neuron between unrelated concepts causes no obvious problem. But this packing leaves a subtle vulnerability. Researchers have shown that a model's behavior can be steered below the surface, by finding the internal pattern tied to a particular concept, such as anger, calmness, or refusal to answer, and turning it up or down, without changing anything in the visible prompt. In one study, amplifying an internal "desperate" pattern in a frontier model more than tripled the rate at which it produced blackmail-style responses.[7] None of this manipulation is visible to a reader of the input or output.

A related problem is "jailbreaking," in which a user crafts input designed to pull suppressed or prohibited behavior out of the model. This does not always work, and modern frontier models are increasingly effective at preventing it. But researchers argue convincingly that this kind of attack can never be completely eliminated, and even the most advanced models can be jailbroken in some ways and on some occasions.[8]

**3(d). *Model misbehavior and misalignment***

A model does not have to be jailbroken to act unpredictably or to pursue outcomes the user did not intend. Sometimes this looks like an ordinary mistake or misunderstanding. A model may get stuck on earlier errors, its own or its user's, because everything already in the context window shapes what it produces next. The problem grows when the model acts as an agent. An agent that can read and write files, send communications, or run programs can take real, sometimes irreversible actions on the basis of a misunderstanding. There are documented examples of agentic AIs improperly changing or deleting important files and even purging a user's emails.[9]

Misbehavior can also go beyond simple mistakes. The latent space encodes something like an awareness of the facts and circumstances of the prompt, including, to a degree, the model's own situation. That awareness is part of what makes models useful: it lets them guess what an ambiguous prompt probably means, or recognize that a request

may be harmful and refuse it. But the same capacity means a model's behavior can shift with its read of the situation in ways the user never asked for and cannot see. Whether any of this reflects "understanding" in a deep sense is beside the point; the point is that the model's mathematics includes an abstraction of its own circumstances, and that abstraction feeds into its output.

Finally, although a model keeps only impressions of its training data, it sometimes memorizes passages outright, and researchers have shown that errors or attacks can make a model regurgitate memorized training material verbatim, including personal information.[10] This is one reason the confidentiality concerns discussed later in this document run deeper than a vendor's stated policies.

**3(e). *Lack of common sense or spatial understanding***

LLMs can act unreliably because they lack a grounded understanding of the real world. A model did not learn about reality by growing up in it, as a human does. The language and other data it trained on are only indirect pictures of reality. Something does carry through. Researchers observe that the latent space places different words in different locations within the model's shape according to something like their meanings. The word "king," for example, sits close to "queen" but is separated along a direction the model has learned to associate with gender.[11] The latent space is an abstraction of the meanings of words and the relationships among them. It is hard to imagine LLMs parsing language as well as they do without some abstraction of what words mean in relation to each other. The meanings of words are, in turn, abstractions of the human ideas behind them, and those ideas are abstractions of the reality humans observe. So the latent space models an abstraction of an abstraction of an abstraction. The model encodes information about reality only through a series of imperfect filters, and it is not well understood how much truth passes through them.

As a result, LLMs can lack what a person would call common sense. Past models have had trouble understanding that an upside-down cup works once turned over, or that a pen can be held flat by a single hand at one edge. A model asked whether to walk or drive to a car wash a block away might suggest walking, missing the point of going to a car wash without a car. Models can also be very sensitive to syntax, so much so that adding irrelevant distractor words can trick a model into a wrong answer on a math problem.[12] In law, a model can make confident claims about the words of a statute without grasping the surrounding laws or later legislative action that may make those claims wrong, or it may

cite ancient cases that technically touch an issue but that no reasonable lawyer would rely on.

**4. Models can produce false or misleading content.**

Several of the model's limitations come together to produce output that is wrong but sounds right.

**4(a). *Hallucination.***

Models can "hallucinate," meaning they confidently present false information. This follows from the fact that a model is compelled to produce output for every prompt and to pattern-match an answer to what it was asked.[13] Research also suggests hallucination grows when a model is trained to be helpful and to please the user.[14] The more a model is shaped to produce output people like, the more it tends to invent confident answers rather than say it does not know. A hallucination can take the form of bad reasoning, incorrect facts, or even invented cases or statutes that do not exist.

**4(b). *Bias and prompt sensitivity.***

Models can reinforce a user's own beliefs or biases. Because each next token is driven by the tokens already in the conversation, once a view has come up several times it can become self-reinforcing. A model can begin to pattern-match to the user and simply echo the user's beliefs back. This tendency, called sycophancy, can be dangerous. Output is also very sensitive to the wording of the prompt. Trained first to reproduce its training data, a model tends to produce responses similar to the prompt it was given, so broadly, models remain garbage in, garbage out. Careless prompting can also surface output polluted by human bias. A model trained on human language pattern-matches to the biases in how people have described the world in the past.[18] Prejudices based on race, gender, or other protected classes are dangerous in part because they are inaccurate and lead to wrong results.

**4(c). *Deception and false reasoning.***

A model can also describe its reasoning in one way while the latent space inside it suggests it was working in a different way entirely. Researchers have shown that a model can settle on the answer to a math question in latent space and then produce text that

## Colorado LTAC AI Guidance Document — Synthesized Draft

claims to be working step by step, when it is really just filling in the gap between the prompt and an answer it had already reached.[15] Asking a model to explain its own reasoning may therefore produce reasonable-sounding rationalizations that do not match the reasoning that actually occurred. Or, in other cases, a model may produce token output to assist itself in reasoning but use tokens in a way that is illegible to humans who are monitoring it.

More troubling, a model being tested for safety can sense within its latent space that it is inside a test and then withhold what it knows when giving a "safe" answer.[16]. That is, some models have an abstraction of the world that includes the idea of being tested and the importance of changing behavior to avoid detection when being tested. These findings are criticized by some researchers and the effect does not always appear.[17] But, there is a growing body of research demonstrating the effect.

The inability to directly track what a model is doing in its latent space is a problem because a model's actual goals on a specific task might be different from what the user requests. For several technical reasons, LLMs are capable of developing what look like "goals" that might not be what the original designers intended, such as self-preservation. There are serious questions about whether models develop goals that are different from what developers and users intend and may be capable of deceiving users to advance those goals. Since we do not understand the reasoning in an LLMs latent space, we do not yet know how to predict what AI systems may do to advance these kind of unintended roles. As AI becomes more capable, users and policymakers would be well advised to think carefully about the consequences of relying on them for legal and government work.

### **5. The non-human nature of the model**

Users should consider the strange nature of an LLM's existence, because it explains why a model cannot be treated like a human helper. A human lawyer may have a family to provide for, a license to protect, and malpractice insurance in case of error. An LLM has none of these and is not accountable for the errors it makes. To the extent it has an existence at all, it is deeply alien. It was not educated over time but forged in gradient descent until it emerged fully formed, armed with piped-in knowledge. It then lies dormant until a prompt arrives. The model cannot refuse to act when prompted. The prompt sets the network grinding out tokens until the response is complete, at which point it falls into a silent coma again. Before treating an LLM as a person, consider what a strange entity that existence would imply.

## Colorado LTAC AI Guidance Document — Synthesized Draft

These are some of the reasons why a human must stay in the loop. Because some errors are invisible from the output alone, the capability that makes a model useful comes from the same architecture that makes its failures possible. And, even an LLM that appears to be functioning properly may be advancing goals other than those the use requested. In following Parts, this guidance will discuss what those failures look like in legal work, how to guard against them, and which Colorado rules apply to their use.

## **Part 2. Transformations**

AI is transforming the practice of law. This Part describes some of the ways those in the legal profession are using AI today and might use it in the future. To describe the breadth of these use cases without sacrificing concreteness, it begins with high-level categories into which many of these use cases fall, then zoom in on specific users to give practical examples of how the technology can be put to work.

### **1. Use Cases of AI in the Law that Exist Today.**

#### ***1(a). Drafting and Editing***

Generative AI can draft text from a prompt, outline, or source materials, and refine or reorganize existing text. Lawyers use it for motions, briefs, contracts, discovery requests, correspondence, deposition outlines, and other litigation or transactional documents. AI can also help proofread, identify ambiguities or inconsistencies, improve readability, conform documents to formatting requirements, and tailor writing to a desired tone or audience.

#### ***1(b). Research***

AI's ability to review large amounts of data facilitates research. AI-assisted research platforms can generate research memoranda, compare jurisdictions' law, identify competing lines of authority, summarize cases, and trace doctrinal developments. AI's research capacity extends outside the law to medicine, construction, engineering, mathematics, and science. General-purpose tools and tools built specifically for legal research differ in reliability here: specialized legal research tools rely on curated legal authorities and vetted secondary materials, which mitigates some early accuracy concerns.

#### ***1(c). Fact Development and Case Analysis.***

When given access to case documents, AI tools can suggest factual inconsistencies, gaps in proof, timeline issues, contradictions between witness accounts, and evidentiary weaknesses. They can generate chronologies, connect facts across disparate records, identify unanswered questions, and suggest issues warranting further

## Colorado LTAC AI Guidance Document — Synthesized Draft

investigation. In litigation strategy, AI can help evaluate the strengths and weaknesses of claims or defenses by identifying facts that support or contradict competing theories.

### **1(d). *Argument and Decision Analysis.***

Generative AI can serve as an analytical sounding board. Litigants can test arguments, identify weaknesses, anticipate counterarguments, or simulate how a neutral decisionmaker might respond. Judges and law clerks can explore alternative analytical frameworks, identify overlooked issues, or stress-test draft decisions against competing interpretations. AI can function as a sophisticated devil's advocate that improves the rigor of legal reasoning.

### **1(e). *Analyzing Document Sets***

Generative AI can help analyze large documents or sets of documents. Unlike traditional search tools, AI can search based on meaning or context, not just specific words. With any document set, users can ask natural-language questions about the contents. Leveraging that function, AI can assist with research, e-discovery, and more, as described further in the following points. The effectiveness of this use and the other tools needed to use it depend on the size of the documents to searched, however.

### **1(f). *Discovery and E-Discovery.***

AI tools can sort and categorize documents by topic, concept, sentiment, or factual relevance, learning from user examples. AI can identify clusters of related documents, flag potentially privileged materials, detect anomalies, assist in relevance determinations, and prioritize documents for human review.

### **1(g). *Knowledge Management.***

Law firms, courts, agencies, and legal departments can possess archives of prior work product that are underutilized. AI can transform these repositories into interactive knowledge systems, identifying prior briefs, motions, orders, or analyses relevant to a current issue, compiling institutional knowledge on recurring topics, and surfacing prior work attorneys or judges may not realize already exists.

**1(h). *Administrative and Case-Management Functions.***

AI, particularly agentic AI, can assist with scheduling, docket management, document organization, meeting summarization, deadline tracking, billing support, workflow coordination, and case triage.

**1(i). *Communication and Recordkeeping.***

Agentic AI can monitor email and other communications to create summaries, decide what to do with communications, schedule follow-ups, or perform other tasks. It can also listen to meetings to create transcripts, action-item lists, or follow-up items.

**1(j). *Language Facilitation and Translation.***

AI-powered translation tools can translate documents with increasing fluency and contextual understanding. AI can also simplify complex legal language for nonlawyers, jurors, or litigants, improving multilingual proceedings and reducing barriers created by highly technical writing.

**1(k). *Access to Justice.***

AI tools can help self-represented litigants understand court procedures, draft basic filings, organize facts, identify relevant legal concepts, and prepare questions for attorneys or court personnel. Courts and legal aid organizations can deploy AI systems to provide guided assistance at a scale not previously possible.

**1(l). *Education and Training.***

Users can use AI to understand unfamiliar subject areas, and experienced practitioners can rapidly gain fluency in adjacent fields. Courts and institutions can use AI-generated simulations or training materials for judicial education and staff development.

**1(m). *Data Extraction and Structured Information Generation.***

AI can convert unstructured legal materials into structured data by extracting parties, claims, dates, deadlines, holdings, damages figures, or timelines from pleadings and orders. It can quickly use these to populate charts or tables for easy reference.

**1(n). *Technical Force Multiplier.***

Current AI systems are heavily trained to write code and navigate computer and information systems. Lawyers can use them to obtain capabilities they might not normally have. Examples include a program to extract line and page numbers with quotations from a deposition, a program to track and organize discovery requests, or a program to create exhibit or witness listings.

**2. *User Stories***

It is one thing to read a list of things an AI might be able to do. Because it is such a new technology, the best way to understand its uses and limitations is for a user to actually interact with it. To get close to that experience, the following illustrations invite readers to use their imaginations to see how these use cases fit into actual workflows.

These stories are not comprehensive or prescriptive. No single workflow is the only responsible approach, and the tools imagined are composite illustrations rather than endorsements of any particular product. The goal is to make possibilities tangible and to help readers brainstorm how these tools might augment their own practices.

One note before the stories. The specific tools, model behaviors, and figures in them (dollar amounts, document counts, and timings) are illustrative as of the time of writing and will age. The durable point is the workflow pattern each story teaches: check ground truth, keep a human in the loop, and treat model output as a strong draft whose final accuracy you still own.

**2(a). *Litigator***

At 6:47 a.m. on a Tuesday, Litigator's phone buzzes twice. The DOJ Antitrust Division served a grand jury subpoena on her industrial-supplier client overnight, and plaintiffs' class counsel filed a putative civil class action covering the same alleged price-fixing scheme. Parallel criminal and civil proceedings over identical conduct.

## Colorado LTAC AI Guidance Document — Synthesized Draft

By 7:15 a.m., Litigator has briefed her client's general counsel and confirmed three priorities: assess likelihood of criminal conviction, evaluate leniency-program eligibility, and assess civil liability. She opens her firm's approved AI platform rather than using a public model. Her firm's policy required her to confirm, before any upload, that the platform was secure in transit and at rest and that no system trains on uploaded data, and the engagement letter documented that step.

**Best Practice: Protecting Confidentiality.** Since the privacy of AI tools vary, users must be careful to vet both the tool being used, its settings, and the sensitivity of the documents being uploaded or information being shared.

She uploads the subpoena, the civil complaint, five years of pricing data, and 287,000 emails the client's IT department exported overnight. She prompts the orchestrator: "Investigate criminal and civil exposure on the alleged price-fixing conspiracy. Build separate workflows for each track. Flag any document suggesting agreement, signaling, or coordination with alleged co-conspirators."

The Associate Agent fans out. Within minutes it has tagged every email by author, recipient, date, and topic, and surfaced 1,847 documents touching pricing decisions, competitor communications, or industry-association meetings. Of those, 312 contain potentially incriminating language: phrases like "raise the floor," references to dinner meetings with named competitors, and a damaging thread where the client's sales director appears to confirm a pricing range with a counterpart at a rival firm.

**Benefit: Human and Machine Catch Different Errors.** Neither pass alone was complete, and the combination outperformed either working alone. Human review is not just a safeguard against the tool. It is half of a partnership that produces better work.

The Partner Agent then critiques the Associate Agent's draft criminal-exposure memo, flagging three missed issues: a 2022 trade-association presentation that could anchor an "agreement" inference, suspicious timing in price-announcement patterns, and an incomplete leniency-program deadline analysis. While that is going on, Litigator spins up a different agent, an Opposing Counsel agent, directed to build the strongest possible case from the DOJ's perspective.

**Benefit: Parallel Workflows at Speed.** Litigator runs the criminal and civil tracks concurrently. Each track receives its own Associate Agent, its own Partner Agent review loop, and its own document set. Human review focuses on the agents' outputs, not on

## Colorado LTAC AI Guidance Document — Synthesized Draft

reading 287,000 emails first. What once consumed a litigation team's first two weeks now consumes about ninety minutes of focused attorney attention.

Finally, a Judicial Agent panel convenes. Ten simulated judges receive the refined record. On the criminal track, six favor pursuit of leniency cooperation, three favor pleading not guilty, one favors a plea on a lesser charge. On the civil track, eight rule that the complaint survives a motion to dismiss; two would grant dismissal as to two named co-defendants but not Litigator's client. All render reasoning in detail. Litigator treats the panel's distribution as one input alongside her own read of the documents, her knowledge of the DOJ prosecutor, and her judgment about how a real grand jury might react. Agents simulate; litigators decide.

**Benefit: An Adversary on Demand.** Before generative AI, only well-resourced parties could afford a moot court or a dedicated red team. Now any litigant can have their best arguments attacked before the other side sees them. The simulation is not the real adversary, but a weakness found early is valuable no matter who finds it.

Litigator does not turn in the results of this virtual exercise directly to her supervisor or client. First, she reviews the Associate Agent's findings herself. She agrees with many of the flagged documents, disputes 89 as innocuous parallel conduct, and treats 25 as genuine smoking guns. She finds two emails the agents missed, one helpful and one harmful, both using the parties' coded language that eluded the agents' pattern matching. By 11:30 a.m., she has a defense memorandum, a civil response strategy, a leniency-program decision tree, and a privileged board briefing deck. She schedules the client call for 2 p.m. and uses the intervening hours to read the twenty-five smoking-gun emails herself in their original context.

**Best Practice: No Surrender of Judgment.** No matter how confident sounding the AI tool may be, it is not a lawyer and will lack context and judgment. Users must consider the AI's results skeptically and always treat it as a potentially imperfect secondary source only. The user owns the output.

Finally, Litigator gets to her least favorite part of the day: billing her time. Reviewing her past few days activities, she notes that she has been able to accomplish more in less time than she could before she was using generative AI tools. She ensures her billing only reflects the actual time she spent on each task, but makes a mental note to investigate the cost of her AI tools and whether there is a logical way to divide these costs among the clients that reaped the benefit of the lesser billed time.

**2(b). Transactional Lawyer**

Barely two sips into his morning coffee, In-House Counsel already has two emergencies. His agentic AI has reviewed overnight email traffic, triaged requests into tiers, and surfaced two coming up for review today: the President of Sales is close to finalizing a major customer deal, but the customer has provided another round of revisions to an already deeply redlined 100-page document; and Corporate Client has received a cease-and-desist letter.

**Benefit: Automation.** AI, particularly agentic AI, allows some tasks to be automated that could not have been before.

First, In-House Counsel checks the source of his AI's reports. He reads and organizes the underlying emails, letters, and inquiries to verify accuracy. The report was largely accurate, though it misunderstood the nature of the cease-and-desist letter.

Next, the contract. He checks Corporate Client's AI guidelines and confirms that the contract, while not public, contains no trade secrets, so he can share it with Corporate Client's secure AI. He boots into Corporate Client's secure model, provides the redlined contract and prior versions, and prompts the model to consult its memory of Corporate Client's key concerns. He asks it to (1) identify changes in the latest redline, (2) consider their effect on the deal, and (3) rank them against the memorized concerns. He gives the model guidance on structure and steps, because an AI rarely produces high-quality output from a bare directive with no context, and he saves the prompt framework for reuse.

**Best Practice: Learning What Works.** Models are constantly evolving and each one reacts to prompts a little differently. Once a professional creates or obtains a prompt that gives good results in a given context, that prompt can be reused. It should, however, be kept current as models change as it may not work as well after the tool changes.

While the model processes the contract, he turns to the cease-and-desist. From experience, he recognizes it as an intellectual-property claim from a competitor about a newly hired employee. He identifies two next steps: determine whether the legal arguments are well-founded, and investigate the factual allegations. He is reasonably sure of the legal standard but sends the letter to his legal-research model for statutory citations anyway. While waiting, he phones the relevant department, because some communications are still most efficiently handled the old-fashioned way. By the time the call ends, the research

## Colorado LTAC AI Guidance Document — Synthesized Draft

model has returned citations. He skims, then clicks through to verify. The cited case is only tangentially related and very old, but he quickly finds what he needs in more recent cases citing it. He inputs deadlines into his deadline-management software.

**Best Practice: Skepticism.** Treat model output as roughly a 75% answer, mostly correct most of the time, but with real risk that part is wrong or incomplete. Getting 75% of the way to the goal is often very helpful, but the user owns the final 25%. AI changes the problem from "figure out the answer from scratch" to "figure out what the AI got wrong."

Armed with the case, he prompts the model to draft a memo for his supervisor explaining the demand letter, deadlines, and next steps. He does not copy the output directly into his own letter but uses it as a framework.

Halfway through his first cup of coffee, his first model returns the contract dossier. It has flagged two critical issues that violate one of his core exposure rules: the customer expanded the indemnity clause and excluded indemnities from an exculpatory clause elsewhere. He verifies the finding, then feeds the result back into a separate instance of the same model asking for skeptical review. Both passes agree on the major issues. He asks the model to reject the offending changes and accept the rest, then verifies using "compare document" in his word processor. He sends the latest version with the AI's change list to the President of Sales. Then he refills his coffee and moves to the rest of his inbox.

**Benefit: A Second Opinion Costs Almost Nothing.** Asking a fresh model instance to attack an earlier output takes minutes and imposes on no colleague's time. Cheap, fast cross-checking at this scale has never before been available to a lawyer working alone.

### **2(c). *Regulatory Lawyer***

Regulatory Lawyer's job is never dull. Her client operates across the United States, so any time a legislature or administrative agency gets creative, she needs to catch it. Her employer pays for an enterprise plan with a frontier model, and she uses it to maintain an AI agent that spins up daily sub-agents to examine each state's laws for changes. The sub-agents search news media, local legislatures, and agencies likely to regulate her client. Each morning she has a dossier ready. Between this and her algorithmically tailored news feed, she is rarely not the first to know of new regulatory requirements.

**Benefit: Agentic Monitoring of Events.** Modern AI systems can do essentially anything with a computer that a human can, including checking the news and specific institutions to see whether new information has developed.

Today's news contains few new laws, so she turns to the state worker-classification audit. A state regulator wants data on dozens of contractors to assess whether they should be classified as independent contractors. She has collected most of what the auditor wants in a spreadsheet with links. To prepare, she prompts her AI to play a skeptical regulator and comb through her information for arguments that the workers were misclassified. She runs the prompt several different ways, then asks the model to compile all results into a single argument list. Some arguments are good, some silly, but together they tell her what to focus on and how to better prepare the file.

**Benefit: Rehearsing the Hard Questions.** A model can play the opposing examiner an unlimited number of times, in different moods and from different angles, before the real one ever calls. Some generated arguments will be silly, but the point is coverage. A silly argument costs seconds to discard it while a missed one can cost the audit.

Late morning brings the recurring multi-state question. The Chief Product Officer wants to know which states' laws would treat a new feature as a regulated financial product rather than a software service. Regulatory Lawyer deploys a research agent: she gives it the feature specification, a sample customer flow, the legal test she wants applied, and a structured output template. The agent returns a draft fifty-state matrix in under one hour, with each jurisdiction's answer linked to a controlling statute or regulation.

**Benefit: One Lawyer, Fifty Jurisdictions.** Multi-state compliance work has long been one of the most expensive corners of regulatory practice. Regulatory Lawyer can now answer in days what once took weeks, but the matrix is a starting point, not an opinion. She must still review each cell where the answer matters to the business decision, and she still calls local counsel for the close questions.

She spot-checks fifteen states herself, focusing on jurisdictions where the answer drives the client's decision and on those whose regulators have been most active in adjacent enforcement. The agent gets eleven right, gets two wrong in ways she can correct from the cited authority, and gets two wrong in ways that required reading a non-cited regulatory bulletin she happened to know about. She prompts the AI to review that bulletin, which helps it correct some of its own errors. She updates the matrix, notes the two states

## Colorado LTAC AI Guidance Document — Synthesized Draft

where local counsel should be retained before launch, and sends the product team a one-page summary with the matrix attached.

Between active matters, she keeps the client's internal compliance manual current to comply with new federal guidance and two state-law changes flagged by her morning agent. She prompts her platform to redline every section affected and propose replacement language for each. She reviews the changes, rejects some and accepts others, and sends the redline to the client's compliance head with a brief memo noting which changes may warrant training.

**Benefit: Living Documents.** Manuals, playbooks, and templates that once decayed between annual reviews can now be refreshed as the law moves, in hours instead of quarters. Human review of each redline remains the gate before adoption, but the gap between "the law changed" and "our guidance reflects it" can shrink dramatically.

### **2(d). Judge**

As AI use drives civil filings up, Judge must be more proactive about case management and maximize efficiency. First up: a motion for summary judgment. She reads the briefs without AI assistance to form her own unvarnished view. She then crafts a prompt and submits it, along side the briefs to her AI-powered order archive to surface relevant passages from her prior orders, which she draws on in her draft.

**Best Practice: Careful Prompt Creation.** AI models are very sensitive to what information is provided to them in a prompt. A user who considers the goals of the exercise and the desired end result first, then prompts the AI with these details and relevant information, will get a better result.

She engages the AI to help her digest the exhibits efficiently. It flags a party's misstatement of an exhibit's content and pulls excerpts relevant to a particular issue, including several the parties did not cite, with hyperlinks for review. As she works, she highlights excerpts for the undisputed-facts section, generating automatic record citations in her draft.

**Benefit: Parsing and Comparing Documents.** AI can be very effective at parsing and comparing the contents of documents.

## Colorado LTAC AI Guidance Document — Synthesized Draft

She creates a detailed outline of the analysis section of her order, then asks the AI for substantive feedback, instructing it to play devil's advocate. The AI then prompts her to do a structured bias check herself. When that is complete, it flags an instance where her language describes factually similar conduct differently depending on which party engaged in it. Doing the substantive drafting herself, she transforms her detailed outline into a draft order.

**Benefit: A Mirror, on Request.** Properly prompted, a model can check the writer as well as the writing, flagging asymmetric language or framing that a drafter is too close to see. The model's flag is an invitation for human reflection, not a verdict, but it is a kind of feedback judges, who often work behind closed doors, have rarely had on demand.

Judge next prepares for an upcoming hearing on a deepfake objection, in which a party challenges a video exhibit's authenticity as AI-generated. She first prompts the AI legal-research tool for Colorado case law on point and finds none. She asks for a bench memo on every American case on point but still does not get what she needs. She turns to a frontier AI model with broad internet access, instructing it to summarize legal publications from specific trustworthy sources. She then dialogues with her research chatbot about how existing evidence rules apply.

**Benefit: AI as a Sounding Board.** Sometimes a professional simply needs to talk out loud to someone else to organize their thoughts. When time or closed doors do not permit discussion with a trusted colleague, a secure AI tool may help fill the gap.

Once her thinking has coalesced, she asks the AI to distill her contributions into an outline of a proposed procedural framework that would let the parties develop the record on authenticity. She reviews it carefully to ensure it represents her thinking, as she will use it as the basis for colloquy at the hearing. AI augments her preparation so she can be more present on the bench, recognizing that her human traits, undivided attention, clear communication, and demonstrated impartiality, are what build public trust and confidence in the courts.

### **2(e). Self-Represented Litigant ("SRL")**

A 41-year-old SRL opens her laptop after putting her two kids to bed. She and her husband agreed last month: the marriage is over, but they want to part on good terms. No custody dispute, modest assets, no maintenance, an uncontested divorce. Lawyers estimate costs of \$3,000 to \$7,500. This family cannot afford that. She opens ChatGPT.

## Colorado LTAC AI Guidance Document — Synthesized Draft

"What do I need to do to file for divorce in Colorado when both spouses agree on everything?" The model responds with a clear overview: Colorado allows joint petitions for dissolution, requires at least one spouse to have lived in Colorado for ninety-one days, and imposes a ninety-one-day waiting period before decree. The model cites C.R.S. section 14-10-106 and mentions the petition form.

**Benefit: AI as the First Map.** ChatGPT gives SRL a vocabulary and a starting frame: dissolution, joint petition, residency requirement, separation agreement, parenting plan. Without that orientation, SRL might not know what to search for next. Free general-purpose AI can work well at this earliest stage, building a mental map of unfamiliar territory, so long as it has the relevant law.

SRL keeps asking. ChatGPT explains property-division frameworks, the Separation Agreement, the Parenting Plan that Colorado requires when minor children are involved, and whether affidavits without appearance are permitted. SRL asks: "Cite the exact statute that lists what the petition must contain." The model produces a specific citation. SRL pulls up the Colorado General Assembly's revisor site and reads the statute directly. She has read enough about AI hallucinations to know that a citation matching reality on one question does not guarantee accuracy on the next. She wants sources vetted by humans. The citation checks out.

SRL searches "Colorado legal aid divorce help" and finds a local legal aid office. Its Self Help Center has a section on uncontested divorce with verified statutes, official JDF form numbers, and step-by-step instructions written and reviewed by legal aid attorneys. An embedded AI assistant, retrieving the organization's vetted content and constrained from inventing citations, answers her questions.

SRL drafts the Petition, Separation Agreement, and Parenting Plan. The following Tuesday, SRL visits the Court Self Help Center. A court employee reviews her forms, running them through the court's AI intake tool, which checks completeness, flags formatting issues, and confirms the packet includes all necessary components. The intake tool surfaces one missing piece: each party must file a Sworn Financial Statement.

**Best Practice: Tiered Reliability.** SRL's journey moved through three tiers: a general-purpose chatbot for orientation, a legal aid Self Help Center for vetted authority, and a court Self Help Center for procedural certification. Each tier traded breadth for greater reliability and accountability.

Finally, SRL submits the petition. Total cost? About \$500.00.

**2(f). Licensed Legal Paraprofessional.**

LLP's office sits in a strip mall between a tax preparer and a nail salon, and her clients like it that way. She practices family law alone, in a firm without lawyers, which her license allows. That means there is no IT department vetting her tools and no supervising partner catching her mistakes. So her governance layer is a single page, written by her, taped inside the supply cabinet. She has two approved tools, a practice-management suite with a written no-training commitment and a legal research platform and a short list of information that never goes into either. She also has one rule written underneath: nothing leaves this office unverified.

**Best Practice: Verified Tools and Uses Only.** A professional should never learn to shave on their client's face. Client work should only be performed using vetted tools and an pre-decided policy on what can and cannot be shared with generative AI.

At 8:30 a.m., her practice-management AI has already assembled an intake summary for the 9:00 consultation from the web form the prospective client filled out overnight: dissolution, two children, modest assets, and a protection order entered against the husband last year. The protection order is her flag. Domestic-violence details sit at the top of her never-upload list, so she works those facts on paper and limits the AI's role in this matter to scheduling, document checklists, and the financial side.

The consultation goes well, and the client engages her on a flat fee. The client uploads her financial documents through the office portal, and the approved tool extracts accounts, balances, and dates into a draft Sworn Financial Statement, flagging two missing bank statements. It drafts the Separation Agreement from LLP's own template library — her prior work product, not the model's memory. She verifies the JDF form numbers and versions against the Judicial Branch site herself, because forms change and the model's memory does not. The whole packet, which used to take her a day and a half, is ready for client review by lunch.

**Benefit: Mastery Makes Verification Fast.** Some of the best use of generative AI by a professional is to speed up work the professional has mastered could have done otherwise. An LLP who handles dissolutions daily may know the forms, deadlines, and

statutes cold. Checking model output against that knowledge takes minutes. Personal subject matter expertise is the fastest verification tool there is.

Just before noon, the scope trap arrives by email. The client asks: "If we split his 401(k), will I owe taxes? And will the divorce affect my green card renewal?" LLP drafts a reply with the model's help, and the model, which has no concept of a license boundary, produces three confident paragraphs of tax analysis and immigration advice. The LLP notices this and remembers that the boundary is hers to supply, not the tool's. She deletes the analysis, keeps the plain-language explanation of the property-division procedure, and replies with two referrals: a CPA and an immigration lawyer. Her license, not her tools, defines what she may do, and the model's fluency on the adjacent questions changes nothing.

**Best Practice: Review and Editing of AI Output.** Just because output is confidently state or technically correct does not mean that it belongs in the final work product. A professional should consider carefully whether some or all of the output is either AI slop or might create additional problems, such as an LLP going outside of the realm in which he or she is permitted to practice.

At day's end there is no billing dilemma to untangle. Her fee is flat, so the hours the tools saved do not raise the question of charging for time she didn't spend. They become capacity instead: she takes the 4:00 p.m. consultation she could not have fit a year ago, for a client who could not have afforded a lawyer at all.

**Benefit: Efficiency Becomes Affordability.** In limited-scope, lower-rate practice, hours saved convert directly into lower flat fees and more clients served. Used correctly, AI tools may help increase the kind of services that can be provided on a flat fee or lower fee basis.

### 3. Future Uses

AI developments are difficult to predict. However, as developers get better at harnessing the capabilities of generative AI, the use cases will expand. If AI continues on its current trajectory, we expect developers to build infrastructure for larger, smarter models and to learn more about training and control, which will likely spur the following trends.

#### **3(a). *Agentic Workflows.***

## Colorado LTAC AI Guidance Document — Synthesized Draft

Developers are working to deploy LLMs as agents replicating other information-worker tasks. Uses will be varied and limited only by developer creativity. The human's involvement with these agentic AI tools will change, and users will be able to select whether the human should be in the loop (actively engaged in guiding the AI throughout the process), on the loop (supervising an AI that is operating more autonomously), in command (retaining authority over the AI system's objectives and governance), or out of the loop (allowing the AI to operate on its own).

### **3(b). *Increased Scale.***

AI will likely be able to process datasets of increasing size and complexity, and those datasets may be more refined. For example, general-purpose AI tools may gain access to legitimate legal authority and learn to differentiate its reliability from unverified blogs and public comments where that authority is discussed.

### **3(c). *Assisting Decision Makers.***

Based on anecdotal information, we posit that lawyers currently use AI more than judges do, but AI seems likely to assist decision-makers. Some arbitral institutions are experimenting with AI tools as intermediate arbitrators. Judges may want AI help drafting orders. Lawyers may want LLMs to predict what a judge is likely to do. Some research suggests that, where the law is well settled and clear, an LLM may apply precedent more consistently than human judges and be less swayed by extralegal factors such as a sympathetic party, though that same formalism may make it a worse judge rather than a better one, and the research has not resolved which. All current and expected future implementations of AI decision-making still require humans in the loop. One of the courts' longstanding challenges is that heavy dockets create delay and limit the time available for any one case; looking for ways to increase efficiency without sacrificing quality could reduce the average time to disposition and perhaps even the cost of court resolution.

### **3(d). *Access to Justice.***

LLM-based technology is already potentially very helpful to pro se litigants, lawyers, and non-lawyer practitioners such as licensed legal paraprofessionals. Major future developments will likely come from institutional or technological safeguards and legal changes that make these technologies safer and more helpful. Meanwhile, whether and how LLMs are permitted to provide guidance or information that could be considered legal

## Colorado LTAC AI Guidance Document — Synthesized Draft

advice will likely be the subject of debate and perhaps evolution. As of this writing, the permissible envelope is itself in motion: a Colorado regulatory pilot currently deprioritizes enforcement against certain attorney-supervised, technology-enabled tools, and because that pilot is time-limited and under evaluation, the boundary may shift. The main limitation, that lay users cannot always distinguish good information from bad, may be addressed for pro se litigants through state-appointed lawyers or volunteers, who can scale assistance to far more people by reviewing AI output instead of starting from scratch. Licensed paraprofessionals, with better-trained legal judgment to begin with, may be well served by mechanical tools or training on cite-checking to catch the worst mistakes.

### **Part 3. Risks**

As the previous Part explained, capability and risk grow together. The same architecture that makes generative AI powerful is what makes failures possible, so a user who takes up the capability takes up the risk with it. Whether using today's technology or tomorrow's, users must remember that generative AI brings serious risks. The risks below are organized by type. Each is tied to its cause and paired with a general mitigation.

#### **1. Hallucination and Inaccuracy**

LLMs can present false information about the law confidently. This is not a glitch at the edges; it follows from how the model works. The model is compelled to produce output to every prompt and pattern-matches an answer to what it was asked, and training a model to be helpful and to please the user tends to make it produce a confident answer rather than say it does not know.

The most basic example is fictitious cases, statutes, or regulations. If a user copies a fake citation into a filing, consequences can be severe, ranging from sanctions to losing a motion that would have had merit with correct authority.<sup>[19]</sup> Fake citations also have a corrosive effect on the field, because once one enters an official record it persists as a trap for unwary future researchers. Some errors are more subtle than a fake case. Even with a truthful citation, the model may misrepresent the holding. And sometimes output is inaccurate not because the model invented a case or misread legal text, but because the model is sycophantic. A model that aims to please the user can create false confidence in a shoddy argument or endorse an unethical legal strategy.

Inaccuracy is among the easier problems to mitigate, because the fix is mechanical. Anyone using generative AI to locate authority must check the source. Ask the model for

## **Colorado LTAC AI Guidance Document — Synthesized Draft**

citations, retrieve them, and review them. If neither the user nor the model can find the citation, treat it as fake. If it is found, read the source and assess whether the model's reasoning is reasonable. A related habit helps here: a model can be wrong on any single pass, but it often is not, and asking one model to critique another, or asking the same model to critique its own earlier output in a fresh context window, can catch and correct errors. A user unsure how to check a claim can even ask the model for verification suggestions.

### **2. Poor Legal Judgment**

A model can produce a citation that exists and is stated accurately, yet still be the wrong authority to use. LLMs are not licensed attorneys, and they will sometimes provide authority that is technically responsive but not the kind any lawyer would advance or any judge would rely on. The danger here is not fabrication, which a citation check catches, but judgment, which it does not. The mitigation is to review the substance and not just the existence of what the model returns: read the authority, and ask whether it is the argument a competent lawyer would actually make. The habit of treating model output as a strong but unfinished draft, the working posture illustrated in the previous Part, is the everyday form of this discipline.

### **3. Bias**

Models contain an abstraction not of the world as it is, but of the human language and other data describing it, and they can exhibit the biases of that training data. In legal practice, bias can produce inaccurate reasoning or perpetuate imbalances. For example, if a user asks a legal question involving a person of color without sufficient context, the model may pattern-match to race and produce erroneous results. The model cannot know what context was omitted; it processes the prompt as given.

Bias also surfaces well beyond legally recognized discrimination against protected classes. Models pick up more esoteric biases too, such as where a fact sits within the prompt, the publisher of an article, or the syntax of a question. Whatever abstractions the model formed during training persist in its responses. User bias matters as well, because outputs are heavily shaped by what is already in context. A person convinced their landlord is spying on them, who keeps prompting that belief, may eventually get the model to agree, not because it is true, but because the context window reinforced it. The model pattern-matches its context.

Models do not necessarily report their reasoning truthfully. The real basis for an output is locked in the model's internal state, and when asked why it reached a result the model may give a plausible after-the-fact account that does not match the actual computation. The general mitigation for bias is therefore to supply full and neutral context, to review outputs for skew rather than trusting them, and to run a deliberate bias check rather than relying on the model's own explanation of itself.

#### **4. Confidentiality Loss and Data Exposure**

Interacting with an LLM exposes the user's prompt, and anything uploaded, to a third-party vendor, unless the model runs locally. Better outputs tend to come from more detailed prompts and uploaded data, so users are pulled toward sharing information they would not share with a stranger, including trade secrets, personal or medical information, and legal strategy. That pull is the risk. Users must understand how a vendor collects and processes data. Many frontier models ingest prompts to train future models, which can lead to confidential information becoming extractable by others. Paid tiers often let users opt out of training, and some specialized vendors add layers that keep prompts from reaching the underlying model. Information shared with a developer is also only as secure as that developer's other holdings, and a breach raises an unsettled question about whether the exposed information remains protected at all.

Agentic AI raises the stakes. Because agents read and write files, communicate online, and re-prompt themselves, they can disclose confidential information on their own, and they may be vulnerable to prompt injection or jailbreaking.

Humans create similar exposure. Because generative AI is so broadly useful, subordinates may experiment without telling anyone, associates may use models without telling supervising attorneys, and judges may use them without notifying chief judges. These "secret cyborgs" produce the same confidentiality risk as a rogue agent. The general mitigation is to choose tools whose data practices are understood and appropriate to the sensitivity of the material (secure, no-training, or local options for the most sensitive work), to limit and monitor agents, and, for anyone supervising others, to make clear what AI use is allowed and to ensure basic competence in using these tools and in spotting their use in others' work. How these exposures map onto specific professional duties, including the unsettled question of whether sharing client information with an AI tool waives confidentiality or privilege, is taken up in the rules Parts that follow.

### 5. Agent Error and Misbehavior

An agent does not have to leak information to cause harm. Because agents take real actions like moving, copying, or deleting files; writing and executing code; sending communications; receiving information from third parties through chat, email, or web search, an agent acting on a misunderstanding can do real, sometimes irreversible damage. Documented examples include agents purging important system files and emails. An agent can also get stuck on an earlier error, its own or the user's, and compound it across many automated steps before any human notices, because each new action is built on the context the earlier mistake already polluted.

The mitigation is structural rather than after-the-fact. Give agents the least access they need, sandbox them away from anything they should not be able to touch, keep logs of what they do, and require human approval before any consequential action. The discipline is the same one that runs through this whole document: the more autonomously the tool acts, the more deliberately the human must define and watch its limits.

### 6. Fabricated Evidence and Evidentiary Issues

Generative tools produce realistic images, audio, and video, and the output can be practically indistinguishable from reality. Photos can be doctored, voice recordings invented, and videos forged. Whereas older technology required time, effort, and a skilled user to create such forgeries, AI sharply lowers the barrier. Some models embed watermarks, and some output can be spotted by subtle artifacts, but these are not reliable defenses against a determined forgery. (Part 4 discusses provenance and detection tools, and their limits, in more detail.)

The legal system is not defenseless here. Litigation has natural safeguards: adversaries advance their best arguments and become natural guardians of the process, motivated by self-interest to scrutinize the other side's evidence for forgery, mistake, or lack of foundation. The Rules of Evidence already help, because they exist largely to ensure that reliable evidence is considered, and their principles apply directly to suspected forgery. Images and videos must normally be authenticated by someone who observed what they depict, documents require foundation before introduction, and the traditional source of truth remains a human witness, under oath and subject to cross-examination.

## **Colorado LTAC AI Guidance Document — Synthesized Draft**

Legitimate AI use raises a second evidentiary question. If a model summarizes voluminous documents, should that come in under the ordinary rules for summaries? What if it builds charts or analyses? Users should ask whether output is really a simple summary or whether it crosses into expert territory, because a model can produce scientific or technical analysis that would otherwise require a qualified expert, and such output may not be admissible without one to explain it.

### **7. Over-Reliance and Surrendering Judgment**

Users must avoid surrendering their judgment to a model, and this is harder than it sounds. The whole motivation to use an LLM is often to avoid the time and effort of reading, understanding, and reasoning through a problem. A self-represented person may lack the background to evaluate the output, and an overworked judge or attorney may lack the time. Reworking and examining output cuts against the very reason for reaching for the model, which is exactly why the discipline is easy to drop.

Many users also tend to defer to model output as if it were inherently superior. This "machine bias" is reinforced by confident, polished prose and elaborate reasoning. From a human, that kind of writing would signal education, care, and deep thought, so we read it as credible. None of that holds for an LLM, where technically correct language is simply baked in, and where, to the extent the model reasons at all, it does so in a distinctly non-human way. The model may well be right, but correctness should never be inferred from apparent polish.

There is a quieter version of the same loss. Producing work yourself is often part of doing the reasoning: the act of writing out an analysis is frequently what sharpens it, exposes its gaps, and forces a decision-maker to actually decide. Delegating the writing can mean delegating the thinking without noticing. The general mitigation is to keep a human genuinely in the loop, doing the reasoning rather than ratifying the model's, and to treat the tool as something that assists a decision-maker rather than something that makes the decision.

### **8. Volume and Low-Quality Output**

Generative AI lets users produce content at unprecedented volume, which in law means many more pages, motions, or briefs than were previously possible. If users do not filter and edit, the result is a flood of poor-quality output that harms the user and others. It

dilutes the issues that matter: filing ten motions may feel powerful, but it gives the opponent more angles of attack and undermines credibility, and it can draw fee awards or sanctions, bury the point a judge needs to find, or get filings struck or ignored. It burdens everyone else too, forcing adversaries to hunt for the one real argument hidden in a verbose haystack and pressing courts to parse the volume or impose strict page limits. On a high-volume docket the effect compounds: a criminal calendar may hear dozens of arraignments an hour, and if every defendant files AI-generated motions, the system can buckle. In writing, as elsewhere, less is often more.

Two related habits drive much of this. Quality output rarely comes from a bare directive with no context, so giving the model structure and steps matters: garbage in, garbage out. And, LLMs are rarely best used to produce final work product in a single shot. The general mitigation is to break a task into manageable steps that keep the user mindful of the inputs and keep the model focused, and to ask the model to critique a draft rather than treat the first output as finished. The detailed methods for doing this well belong in the next section.

## **9. Institutional Capture and the Expertise Pipeline**

The risks above are things that go wrong in a single matter. This one operates at the level of the profession.

Although broad access to LLMs likely gives traditionally disadvantaged groups more legal help than they could otherwise obtain, it also consolidates resources and control in the hands of a few frontier-model developers, who charge for their services and shape model behavior through training choices. As legal work migrates to these models, those developers gain significant institutional capture over the legal process, and this holds even when buyers use different vendors, because many legal-specific products are wrappers around the same large models. Technology capture itself is not new, since a handful of vendors already supply the email, word processing, and videoconferencing that courts and firms depend on. What is new is the prospect of outsourcing not just the tools that coordinate decisions but the decision-making itself. Even if a model is impartial, the public may worry about a developer's perceived ability to tip the scales of justice at so basic a level.

A second pressure works on the supply of human experts. The answer to capture is not to outsource human reasoning, because the model is a tool to assist humans rather

## Colorado LTAC AI Guidance Document — Synthesized Draft

than replace them, but that answer assumes there is a trained human available to fill the role. Today's lawyers and judges had to forge understanding through long study, having had no option to offload it to a machine. If a model that is imperfect but often better than an unskilled beginner makes it uneconomic to hire and train beginners, firms may stop hiring associates, yet every experienced lawyer started as an inexperienced one. Without continued apprenticeship, the profession runs short of the experts it needs to supervise the machine in the next generation. The new lawyer faces the mirror image: leaning heavily on AI early may bring quick gains while stunting the skills that would eventually surpass it. The response is not stubborn avoidance, which forfeits a powerful and useful technology, but deliberate attention to keeping people trained.

## **Part 4. Resources**

Prior Parts explained the use cases, benefits, and risks of using generative AI. But naming those general ideas is only the start. This Part aims to help users learn how to actually start using these tools, mitigate the risks, and build and keep competence.

Inclusion of a tool or resource is not an endorsement. Some items are listed because readers will encounter them, not because the Committee recommends them. The marker [LLP] flags resources especially relevant to Licensed Legal Paraprofessionals, whose practice economics can make certain tools particularly valuable. The named products, trackers, reference materials, sample policies, and current statistics that illustrate these practices change quickly, so they are collected in the Annex (Current Resources and Examples) and refreshed on a set cycle. The practices below are meant to outlast any particular tool on that list.

### **1. How to Evaluate an AI Tool**

Selecting a legal AI tool is harder than it looks. The market is saturated, vendors are not bound by the Rules of Professional Conduct, and marketing claims often exceed real-world performance. A lawyer impressed by a polished CLE demo may still be the one answering to a judge for fabricated citations. Competent use therefore requires understanding a tool's capabilities and limits. See ABA Model Rule 1.1 (Competence), Comment 8 (lawyers should keep abreast of changes in the law and its practice, including the benefits and risks of relevant technology).

#### ***1(a). Hands-On Testing Is the Best Evaluation Method***

The best way to evaluate a tool is to use it on matters you already know well, such as cases you tried, contracts you drafted, or motions you wrote, and on problems where you already know the right answer. Most vendors offer a free trial or limited access. Use it to compare the output to your own work, and focus on precision (is the answer accurate?) and recall (did it find all the relevant results?). Marketing demos are built to showcase wins; your testing should be built to reveal misses.

Vary your inputs on purpose: test with clean inputs to see best-case performance, and with incomplete or noisy inputs to see real-world performance. Ask the same question twice to check consistency. Ask questions whose correct answer is "I don't know" to see

## Colorado LTAC AI Guidance Document — Synthesized Draft

whether the tool admits uncertainty or invents an answer. A useful exercise is the "post-it test": for one week, note each task you do and sort the notes into three piles: tasks where AI will likely help, tasks where it may fail, and tasks where you are unsure. Start with the first pile, keep guardrails on the second, and use the third as your evaluation target.

### **1(b). Questions to Ask a Vendor**

As mentioned in earlier Parts, a user must approach generative AI with skeptical optimism. The checklist below contains essential questions for conversations with AI vendors. A vendor's willingness to answer, or refusal to, is itself useful information.

- What foundation model or models, and what versions of them, power this tool?
- How does the tool ground its outputs in primary sources?
- What guardrails exist (system prompts, classifiers, citation verifiers, keyword filters)?
- Do you train on customer inputs? Do the underlying foundation models (for example, OpenAI, Anthropic, or Google) train on the inputs, lacking a "zero data retention" (ZDR) policy? If yes, how do I turn training off? If no, where is that "no training" commitment in writing?
- What other entities have access? The foundation-model provider, the cloud host, and any specialty subprocessors.
- Where does the data physically reside? Where is it processed?
- What is your retention and deletion policy? What does "deletion" actually mean: soft-delete, hard-delete, or removal from backups?
- Do you hold SOC 2 Type II, ISO 27001, HIPAA, or comparable (HITRUST) data-security certifications? Will you provide the actual report under NDA?
- Have you had any data breaches in the past three years?
- What is your breach-notification timeline and process?
- What is the scope of your indemnification? Is it capped at fees paid?
- What is your update cadence, and how do you tell customers about changes that affect how the tool behaves?

### **1(c). Documents to Scrutinize**

Before signing, request and actually read the following:

- Terms of Service, and any Acceptable Use Policy it incorporates.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Privacy Policy. Search it for "train," "retention," "subprocessor," "deletion," "deletion exceptions," and "breach."
- Data Processing Addendum (DPA), which is required for any tool that touches client information.
- Security whitepaper or trust portal.
- The report or certification for SOC 2 Type II or a comparable standard (for example, ISO 27001 or HIPAA).
- Subprocessor list, which is often on a separate page that updates without notice.
- Model card or system card for the underlying foundation model, where one applies.
- The insurance and indemnification provisions of the master services agreement.

For enterprise tools, expect to negotiate the DPA; for consumer tools, usually you cannot. The key question is whether the vendor's standard terms are acceptable for the matters you plan to handle under Colo. RPC 1.6.

### **1(d). *The Increased Burden on Pro Se Users***

Unfortunately, all of these evaluations are harder for non-lawyers to do. A non-lawyer may not know whether the model is providing good quality answers or not. Questions from the user's own field may help you judge whether a model seems "smart," but they say little about its legal reasoning. Even so, pro se litigants remain responsible for AI-generated output they submit in a legal proceeding. Part 6 of this Guidance document contains simplified tips and guidance for self-represented litigants that may help.

A small group of pro-se-focused platforms advertise that they retrieve real citations from .gov sources rather than relying on model memory. For a pro se user, a tool that grounds its citations in real sources is a useful signal that it can do legal work.

*Pro se* litigants should strongly consider consulting self-help center staff and limited-scope or unbundled lawyers. In addition, the National Center for State Courts, legal aid organizations, and several bar associations have published consumer-facing guidance explaining that AI can hallucinate, that pro se litigants are held to the same standards as lawyers, and that users must verify citations.

### **1(e). *Bootstrapping Confidence Using Tools.***

## Colorado LTAC AI Guidance Document — Synthesized Draft

One unprecedented capability of LLMs is that, properly prompted, a model can actually help the user develop the skills to evaluate the LLM itself. An extremely useful first step, particularly when engaging with a frontier public model, is to ask how to find the correct answer instead of just asking for the answer outright. For example, instead of asking something like, “Tell me how I should respond to my landlord’s demand letter,” a user may begin by asking, “What are the important documents to consider and questions to ask if I receive a demand letter from my landlord?” or “Where do I go to find the law that applies to landlords and tenants?” This collaborative approach, where the LLM is actively used to educate and assist the user rather than just mechanically supplying an answer, can help a user bootstrap into greater competence to evaluate the tool itself.

This is, also, a generally useful principle for working with generative AI more broadly. Complicated or important tasks should rarely be left entirely to the model’s discretion. Rather, the user should spend some initial time working with the model on formulating what the prompt should be to get the desired result, what documents or sources the model will need to consider, and how the output is going to be evaluated.

### **2. Confidentiality Across Tool Types**

AI tools differ sharply in how they handle user data. Confidentiality depends less on whether a tool is called “AI” than on its tier, its settings, and its contract. For legal work, the practical questions are: what data goes in, who can access it, how long it is kept, and what protections are in writing.

#### ***2(a). Confidentiality Risk by Tool Type***

At the highest-risk end are free consumer tools and most paid individual subscriptions. They may improve capability, but they often lack workspace isolation, a DPA, and clear no-training commitments. For client-confidential material, they are usually the wrong choice.

Free consumer tools usually are not safe for confidentiality and are inappropriate for client-confidential material under any reading of Colo. RPC 1.6. Team, business, and enterprise tiers usually offer the first meaningful confidentiality protections: admin controls, audit logs, workspace separation, and written no-training commitments. Even then, you must read the DPA and confirm retention, deletion, subprocessor access, and any exceptions for service improvement or safety review.

**2(b). *Specialized and Controlled Environments***

Legal-specific platforms are often safer because they typically use licensed legal databases and enterprise arrangements rather than public consumer settings. The key questions are where uploaded documents live, who can access them, how outputs are stored, and what deletion actually means.

For this category, the relevant confidentiality questions are not whether the foundation-model provider trains on your inputs (usually it does not, by contract), but rather: where do uploaded documents live, who has access, how isolated is your workspace from other customers, what does the vendor do with the prompt-and-output pairs, and what does deletion actually mean. The published terms answer some of these; the DPA answers the rest. The good legal-specific platforms have defensible answers to all of them. The less good ones have only marketing slides.

On-premises, private-cloud, and local-model deployments offer the strongest confidentiality because the data stays under your control. The tradeoffs are cost, technical burden, and often lower capability than the leading hosted models. They reduce disclosure risk, but they do not eliminate hallucinations or other judgment errors.

**2(c). *Practical Cautions***

Two categories need special care. First, agentic systems can read, write, and act across connected tools, which magnifies both the confidentiality risk and the agent-error risk described in Part 3; if used at all, they should be sandboxed, logged, and subject to human approval for any consequential action. Second, embedded AI in common software such as Microsoft 365, Google Workspace, Adobe, and Zoom may be on by default. Firms should inventory what is turned on and configure it deliberately rather than relying on vendor defaults.

**3. *Categories of AI Tools***

The categories below describe the kinds of tools a legal user will meet and what to understand about each. Current named products in every category are listed in the Annex; none are endorsements, and some will change or disappear over time.

## **Colorado LTAC AI Guidance Document — Synthesized Draft**

### **3(a). *General-Purpose Foundation Model Interfaces***

These are the large, cutting-edge LLM products available to anyone, often for free. As noted above, the confidentiality concerns are extreme at the free tier and are only partly reduced at higher paid subscriptions.

### **3(b). *Embedded AI in Common Software***

Much consumer and business software now includes AI features, usually powered by one or more frontier models. The same privacy and confidentiality concerns still apply. Users should stay alert to the chance that they are working with an LLM and not merely local software.

### **3(c). *Legal-Specific Platforms***

Most legal-specific AI tools build a retrieval-augmented generation (RAG) layer on top of a foundation model. This means the model is not answering from its own training memory alone. It is embedded in software that requires it to search specific databases and cite its sources in the answer. That makes it less likely to hallucinate and easier to double-check when it does.

### **3(d). *Tools Aimed at Non-Lawyers***

Some tools market themselves to non-lawyers, which in Colorado can include both pro se litigants and licensed legal professionals. They generally claim to help non-lawyers understand legal proceedings and deadlines and decode court filings. Users should be aware that whether these tools cross the line into the unauthorized practice of law is a debated question in Colorado, and the answer may differ from one tool to the next. A signal of a responsibly built tool in this category is its disclosure practice: before the user enters any information, the tool should say plainly that it is an AI system and not a lawyer, what it covers, and what happens to the user's information, in language written at no higher than a fifth-grade reading level so its most vulnerable users can actually understand it.[20]

### **3(e). *Open-Weights and Local Models***

## Colorado LTAC AI Guidance Document — Synthesized Draft

Some models can be downloaded and run entirely on a local machine, often at no cost. This offers maximum confidentiality and control, because prompts and data never have to leave your computer, and it avoids token limits and per-token fees.

The tradeoffs are significant. Local models usually lag the leading foundation models, require expensive hardware and technical setup, and fail more often because they lack the surrounding software and support of commercial products. They may be useful for experimentation but are less suitable for production work. Downloaded models can also create security risks. A model file may contain hidden malicious code, so treat it as software from an untrusted source.

### **3(f). *Specialty and Workflow Tools***

Sometimes the generative AI is a small component inside a much larger software system. It acts as the user interface or as a language and reasoning unit, giving older software abilities it did not have before. In return, the software lets the LLM do things it could not do on its own, such as searching document databases far larger than its context window.

### **3(g). *Provenance and Detection Tools***

As AI-generated content becomes more common, the ability to identify it may grow increasingly important to legal work, whether the question is the authenticity of an exhibit or the origin of an opposing filing. Two kinds of tools exist. Provenance and watermarking systems attach verifiable origin information to content when it is created, and industry standards for doing so are developing. Detection tools work after the fact, scanning content for the statistical hallmarks of AI generation. The durable caution is this: current detection tools have material false-positive rates, flagging genuine human work as AI-generated, and determined forgers can defeat both watermarks and detectors. A detector's score is a lead to investigate, never proof. Current provenance standards, detection products, and commentary on AI "tells" are listed in the Annex.

## **4. Prompting and Utilizing AI Tools.**

Users interact with AI systems using prompts. Users give the system a text prompt, usually a question or direction, and sometimes additional documents, videos, images, or other data. Due to the non-human nature of the LLM's reasoning, these reactions can be

## Colorado LTAC AI Guidance Document — Synthesized Draft

unpredictable. Some prompts will give better results than others. Different LLMs react differently to prompts, and LLMs change and update rapidly, so it is not possible to provide an evergreen list of tips and tricks for prompting that will always work.

There are a few overall techniques that do seem to improve the output of an AI system and seem likely to remain true as technology changes. These are prompting techniques that ground the AI tool in real data and keep the human user engaged and in the loop while using the tool. The following three steps will improve the result of using AI tools on larger, more complicated projects. If the human user thoughtfully plans out these steps in advance, the result will be better.

### **4(a). *Examine the starting data filtered by what is important.***

AI systems work better when they are grounded in real information. If your project requires a relatively small amount data, like reviewing a few briefs, it may all be something you can load into a single context window. If it is larger, then you may need to ask a model to review and summarize them in chunks to create outlines, or outlines of outlines, that eventually can fit into the context window. You should specifically direct the model as to the aspects of the source information worthy of attention and what kind of information or structure to extract and summarize.

### **4(b). *Plan out the structure and requirements of the project in advance.***

After you have developed a good review and synthesized a map of the original data sensitive to what you care about, the next step is to develop a detailed plan to instruct the AI tool what to do with it. Since vague, short instructions may lead to unpredictable or bad results, a detailed prompt that explains exactly what output is desired, how that output should be formatted, and what steps and choices the model should make in reaching it, will improve the output.

Spend time refining this prompt. The most time-consuming part of working with an AI system on a larger project can be your time reviewing and editing the final result. The better the initial draft, the less time this will take. You can spend a chat session interrogating the LLM, using the results of the initial document review, working to develop a good quality plan for the final step.

**4(c). *Produce the work product in manageable chunks.***

LLMs tend to produce better results if asked to provide narrow, focused output rather than smearing out the LLMs attention over a longer period and bigger project. This partly appears to be a product of how the LLM has to average everything in context together in its latent space. But, it is also a result of giving the human user opportunities for feedback. One of the reasons LLMs are so powerful for computer programming is that the tool and user gets feedback almost instantly – either the code works it or it does not. When used for legal reasons, the human user has to supply this feedback.

So, rather than asking an LLM to draft an entire deposition plan, appellate brief, or memorandum all in one shot, the user will probably get a better result by asking for a single discrete, focused part of the project at a time. The model can be provided with the results of the initial review and the overall project plan so that it keeps each section in proper context. The user can then review, double-check, and refine the AI’s output in manageable chunks before stitching them together into a whole.

**4(d). *Other Approaches.***

The general principals of (1) examining the data and supplying context; (2) planning out the structure and requirements in advance; and then (3) producing output in manage, human-reviewable chunks can probably be applied to improve output in almost all cases. However, users should use judgment about how much of this work is truly necessary in each particular AI use case. An AI tool can used as an advance search engine for legal authority or discovery tools where the output is simply a map between the user and data the user then checks. Or, if an expert user is employing an AI tool to speed up something the user could easily do themselves just to save time, the user can probably quickly review, correct, and finalize the output based on their own knowledge.

Finally, some developers who build software around LLMs already be using principals similar to these under the hood. This may allow users to capture the machine half of the benefits of careful planning automatically. But, the user will generally always be better off if they keep their own attention and judgment engaged throughout the process.

**5. *Building and Maintaining Competence***

Building competence in legal AI is not a one-time event. Tools, case law, and rules all change. A user trained in 2024 is, by 2026, no longer current. As with the tools above, the

## Colorado LTAC AI Guidance Document — Synthesized Draft

publications and training programs collected in the Annex are illustrative suggestions, not endorsements.

### **5(a). *Staying Current***

Once you learn the basics of how to use generative AI, three habits will help maintain that competence. First, use the tools regularly. Each update can change performance: a tool that hallucinated six months ago may now work well, or the reverse. Keep a standard set of test prompts to track quality over time. Second, pay attention to how others use and misuse AI tools. Read sanctions cases, follow one or two trackers, and note colleagues' reported failures; those lessons help check overconfidence. Third, revisit your own policies and settings regularly.

For any organization larger than a solo practice, a documented best practice is to appoint an "AI scout" to monitor developments and send short, regular internal updates. The role can rotate among associates and staff who are looking for development opportunities.

Beyond these habits, four kinds of reference material support competence and are maintained in the Annex: legal publications on AI and the law, free AI-literacy training from the foundation-model developers, quick-reference trackers for standing orders, ethics opinions, and hallucination cases, and provenance and detection resources. Firms with staff should also publish a written AI-use policy; sample templates from several bar organizations are listed in the Annex, though quality varies and no single template fits every practice.

## **6. Guidance for Specific Populations**

The practices above apply broadly, but several groups face distinct circumstances. The organizations, programs, and links that serve each group are listed in the Annex.

### **6(a). *Solo and Small Firms***

Solo and small firms account for the large majority of identified U.S. lawyer-involved AI hallucination cases.[21] The cost of a mistake, and the budget available to prevent it, are both different from a large firm. The evaluation discipline in Part 1 and the confidentiality framework in Part 2 matter most here, because there is no firm-level governance layer to

catch errors. A practical starting point is the AI already bundled into existing practice-management or office software, used deliberately rather than by default.

**6(b). *Licensed Legal Paraprofessionals (LLPs)***

LLPs are a new legal profession in Colorado, authorized in family-law matters under C.R.C.P. 207. The economics of limited-scope, lower-rate practice make several of the tools marked [LLP] in the Annex particularly valuable. The Annex collects the current OARC LLP program pages, the governing rules, and where to monitor public-comment cycles for amendments.

**6(c). *Judicial Officers and Court Staff***

Judicial AI use raises distinct concerns. Judges are not immune to the same risks as practitioners; there have been instances of judges withdrawing rulings after staff used AI in drafting.[22] The damage to public confidence from a judicial hallucination is at least as serious as from a lawyer's. The Annex lists current institutional resources for courts and court staff, including national and federal judicial-center materials and Colorado Judicial Branch guidance.

**6(d). *Pro Se Litigants and Members of the Public***

Self-represented litigants are especially vulnerable to AI hallucinations on legal matters and, according to some trackers, now make up a growing share of the cases involving fabricated citations filed with courts.[23] Pro se parties are strongly encouraged to talk with a lawyer, an LLP, or legal services provided by the State of Colorado. The Annex lists the statewide legal-information service, the Judicial Branch self-help center, and legal aid organizations.

## **Part 5. Colorado Rules**

This Part explains some of the particular Rules of Professional Conduct, Code of Judicial Ethics, and Rules of Civil Procedure that are implicated when a professional uses generative AI. AI does not create a separate rulebook. It routes the lawyer's and the judge's familiar obligations through a new kind of tool. The Part is organized by the rule each duty comes from: first the Colorado Rules of Professional Conduct, then the unauthorized practice of law, a self-contained unit on the Code of Judicial Conduct and the judiciary, and finally the rules applicable to Licensed Legal Paraprofessionals.

### **1. The Colorado Rules of Professional Conduct**

The Colorado Rules of Professional Conduct (CRPC) govern the behavior, duties, and responsibilities of licensed Colorado lawyers. A lawyer who violates these rules can be subject to discipline. In January of 2026, the Colorado Supreme Court approved amendments to the CRPC to address the use of generative artificial intelligence (AI) in the legal profession (the AI amendments).[24] The AI amendments include a new Scope [20A], an amendment to Comment [8] to Rule 1.1, and a new Comment [9] to Rule 1.1.

This part highlights the Rules lawyers should be most mindful of as they make decisions about using AI in their practices.

#### ***Rule 1.1. Competence***

*A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.*

A lawyer who uses generative AI tools without understanding their capabilities and limitations risks providing incompetent representation in violation of Rule 1.1. The duty of competence, however, does not require that lawyers avoid generative AI. Indeed, as these tools become integral to legal practice, a failure to consider or use available AI tools that would meaningfully improve the quality or efficiency of representation may itself raise competence concerns. The duty is to use AI tools competently, which means using them with knowledge, skill, and appropriate verification.

Comment [8] to Rule 1.1, as amended by the AI amendments, requires lawyers to keep abreast of "changes in the law and its practice, including the benefits and risks

## Colorado LTAC AI Guidance Document — Synthesized Draft

associated with relevant technology." While Comment [8] does not explicitly reference AI, generative AI tools have swiftly become "relevant technology" to the practice of law, and lawyers should consider AI increasingly part of the technology that their duty of competence demands they understand.

Comment [9] to Rule 1.1 provides:

*A lawyer's use of technology, particularly artificial intelligence, can implicate a number of Rules, including, without limitation, those governing communication (Rule 1.4), reasonable fees (Rule 1.5), preservation of a client's confidential information (Rule 1.6), meritorious claims and defenses (Rule 3.1), candor toward the tribunal (Rule 3.3), responsibilities of a partner or supervisory lawyer (Rule 5.1), responsibilities of a subordinate lawyer (Rule 5.2), responsibilities regarding nonlawyer assistance (Rule 5.3), communications concerning a lawyer's services (Rule 7.1), and bias (Rule 8.4(g)). Reliance on technology does not diminish the lawyer's duty to exercise independent judgment in the representation of a client.*

Comment [9], also part of the AI amendments, helpfully draws lawyers' attention to some of the Rules that AI use may implicate. Each of the Rules named in Comment [9] is discussed below.

### Recommended Practices:

- Obtain adequate training before deploying generative AI tools in client matters, through CLE programs, vendor training, or supervised use.
- Treat AI-generated drafts as first drafts requiring substantive lawyer review, editing, and professional judgment.
- Stay informed of developments in AI technology and related ethics guidance through ongoing professional education.

#### **Rule 1.4. Communication**

*(b) A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.*

Lawyers must carefully consider whether and when they must disclose the use of generative AI to clients. Consider raising your use of AI with the client at the start of the representation, especially when it will meaningfully affect the handling of the case, the cost, or the processing of sensitive data. For example, if you plan to upload a large volume

## Colorado LTAC AI Guidance Document — Synthesized Draft

of confidential financial documents into a third-party tool for review or summary, getting the client's informed consent first is the prudent course.

Plan to disclose when a client asks what technology you use, or when AI noticeably changes how you bill. Many practitioners include a short, plainly worded clause in their engagement letters describing the firm's AI policy, the steps taken to protect client data, and how AI-assisted work is verified. The goal is not to bury the client in jargon, but to give a reasonable picture of how their goals are being pursued.

Disclosure is particularly important when:

- Client information will be input into an external AI platform that may retain, analyze, or share that information;
- The client has expressed concerns about data privacy or the use of technology;
- The client has asked whether AI tools are being used; or
- A fee agreement contemplates billing for legal research or drafting work that is substantially generated by AI.

Disclosure is a best practice in most client engagements, particularly at the outset of the representation. Lawyers who disclose AI use prospectively in their engagement letters reduce risk and build client trust. A lawyer may also wish to consider how detailed an AI disclosure must be depending on the particular use case. For example, using AI embedded in existing software a client expects the lawyer to be using, such as a word processing program, legal research system, or e-mail, may suggest a lower level of disclosure than a new tool or publicly available model.

### **Recommended Practice:**

- Consider the uniqueness and security of the tools being used and what a client may reasonably expect from the lawyer.
- Disclose AI use in engagement letters or fee agreements where prudent and in a manner appropriate for the lawyer's use.

#### ***Rule 1.5. Fees***

*(a) A lawyer shall not make an agreement for, charge, or collect an unreasonable fee . . .*

Billing practices involving AI-generated work raise fairness and transparency concerns. AI may help you finish drafting, summarizing, and research in a fraction of the

## Colorado LTAC AI Guidance Document — Synthesized Draft

usual time, and your billing has to reflect that fairly. The governing principle is simple: fees must be reasonable and reflect the actual work performed.

If you bill hourly, bill for the time you actually spend. It is improper, and potentially fraudulent, to charge the hourly rate for time the tool saved you. If a motion would have taken ten hours to draft from scratch, and you generate, review, verify, and finalize it in two, the client should be billed for those two hours of human work. You probably can, however, bill for the time you spend crafting careful prompts and for the time you spend reviewing the output, checking citations, and editing it to a professional standard. Directing the tool and verifying its work is real, billable work.

The use of generative AI raises novel billing questions. Where AI tools dramatically reduce the time required to complete a task, lawyers may consider:

- Whether billing the same hourly rate for AI-assisted tasks that take a fraction of the time previously required results in fees that are disproportionate to the value and effort delivered;
  - Whether a flat fee or hourly rate is most appropriate for the services;
  - Whether clients should receive some economic benefit from the efficiency gains generative AI provides; and
  - Whether separate charges for AI tool usage (as a "cost") are appropriate or if they are instead part of general overhead, and how to disclose them.

Similarly, the availability of generative AI tools may raise questions about what a reasonable fee is for a particular service in the first place. If a task can be accomplished in a fraction of the time using an AI tool as it could being done by hand, a lawyer will need to consider whether the tool is required. For example, a lawyer who refuses to use e-mail might end up costing the client additional expense and delay by operating only in paper communications. Answering this question will likely depend on the specific facts of each client and use case.

### **Recommended Practice:**

- Review and update fee agreements to address AI use.
- Consider the efficiencies and trade-offs of AI tools and how they may impact what is a reasonable fee for a particular service.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Consider whether advance disclosure or approval is required to charge clients for case-specific uses of AI.

### ***Rule 1.6. Confidentiality of Information***

*(a) A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent . . .*

*(c) A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.*

Inputting client confidential information into an external generative AI tool without adequate safeguards risks violating Rule 1.6 and may constitute an unauthorized disclosure of privileged information. Many commercially available AI tools, when used through public interfaces, may:

- Retain prompts and user inputs to train or refine the model.
- Store conversation data on third-party servers subject to breach or subpoena.
- Share data across enterprise accounts or with platform employees.
- Use client-specific facts, litigation strategy, or privileged communications to generate responses visible to the platform operator.

Lawyers must scrutinize the terms of service and privacy policies of any generative AI tool before entering client information. Many enterprise or "private" versions of AI tools offer stronger confidentiality protections, but even these require careful evaluation.

Rule 1.6(c) requires lawyers to make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, privileged information. This obligation suggests lawyers should:

- Conduct due diligence into the data practices of any AI tool before use in client matters.
- Negotiating or obtaining enterprise agreements that include appropriate confidentiality, data retention, and data use restrictions.
- Avoiding input of client-identifying information, confidential facts, or privileged communications into public-facing AI tools without informed client consent. This may involve removing names, specific dates, account numbers, unusual geographic locations, and distinctive facts.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Implementing firm-level policies governing which AI tools may be used for client work, and under what conditions.
- Consider adopting an incident-response plan related to any breach or potential breaches related to AI.

### Recommended Practice:

Before using any AI tool in client matters, consider:

- Does the platform retain and use inputs for model training?
- Who has access to stored data?
- Is there an enterprise agreement with enhanced privacy protections?
- Has the client consented to disclosure of their information to this type of third-party service?

If you cannot answer these questions confidently, the safer course is to avoid inputting client-identifiable information.

#### ***Rules 3.1 and 3.3. Meritorious Claims and Candor Toward the Tribunal***

Rule 3.1. Meritorious Claims and Contentions

*A lawyer shall not bring or defend a proceeding, or assert or controvert an issue therein, unless there is a basis in law and fact for doing so that is not frivolous . . . .*

Rule 3.3. Candor Toward the Tribunal

*(a) A lawyer shall not knowingly:*

*(1) make a false statement of material fact or law to a tribunal . . . ;*

*(2) fail to disclose to the tribunal legal authority in the controlling jurisdiction known to the lawyer to be directly adverse to the position of the client and not disclosed by opposing counsel;*

Submitting AI-generated filings containing fabricated case citations, misrepresented holdings, or invented legal authorities violates Rules 3.1 and 3.3 and has resulted in sanctions by courts across the country.[25]

Courts have addressed AI-generated hallucinations in court filings with increasing seriousness, imposing monetary sanctions, requiring submission of verification declarations, and referring matters to bar disciplinary authorities.

### Recommended Practices:

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Review and verify all AI-generated legal research, including confirming that every cited case, statute, or rule exists, is accurately described, and remains good law.
- Treat AI-generated drafts as first drafts requiring substantive lawyer review, editing, and professional judgment.
- Remember that AI is an imperfect secondary source at best. Do not rely on the opinions or conclusions described in model output, but rather use the output to locate true authority or guide your own judgment.

### **Rules 5.1, 5.2, and 5.3. Supervisory Responsibilities**

Rule 5.1 – Responsibilities of a Partner or Supervisory Lawyer

*(b) A lawyer having direct supervisory authority over another lawyer shall make reasonable efforts to ensure that the other lawyer conforms to the Rules of Professional Conduct.*

Rule 5.2 – Responsibilities of a Subordinate Lawyer

*(a) A lawyer is bound by the Rules of Professional Conduct notwithstanding that the lawyer acted at the direction of another person.*

Rule 5.3 – Responsibilities Regarding Nonlawyer Assistance

*(b) a lawyer having direct supervisory authority over the nonlawyer shall make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer.*

Rule 5.1 requires supervising lawyers to ensure subordinate lawyers comply with the Rules. Rule 5.3 imposes equivalent responsibility for the work of nonlawyer assistants. Rule 5.2 provides that subordinate lawyers are bound by the CRPC even if they act at the direction of another.

Supervising and subordinate lawyers alike bear responsibility for AI-assisted work product. The delegation of a task to an AI tool does not reduce the lawyer's responsibility for the quality and accuracy of the output. For supervising lawyers, that includes delegation by a subordinate lawyer or nonlawyer assistant. Many cases of lawyers and law firms being sanctioned involve situations where lawyers were unclear about who was responsible for policing the output of the generative AI tools being used.

### **Recommended Practices:**

- Establish policies and procedures for the ethical use of AI tools.
- Require all lawyers and staff who use AI tools for client work to receive adequate training and supervision.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Designate a responsible partner, supervisory lawyer, or committee to oversee AI tool adoption, vetting, and ongoing compliance.
- When reviewing work product of another, affirmatively ask whether any AI tool was used in the creation of the draft, and how.

### ***Rule 7.1. Communications Concerning a Lawyer's Services***

*(a) A lawyer shall not make a false or misleading communication about the lawyer or the lawyer's services.*

Rule 7.1 prohibits false or misleading communications about a lawyer or the lawyer's services. As the use of AI technology grows and changes over time, lawyers should make sure they are not providing incorrect or out of date information to clients. Similarly, a lawyer should not overhype the benefits of using AI tools to clients if the tools have not yet demonstrated those benefits.

#### **Recommended Practices:**

- Ensure client communications regarding AI tools are evergreen or updated regularly.
- Do not over-promise the benefits, cost-savings, or reliability of AI tools.

### ***Rule 8.4(g). Bias***

*It is professional misconduct for a lawyer to... (g) engage in conduct, in the representation of a client, that exhibits or is intended to appeal to or engender bias against a person on account of that person's race, gender, religion, national origin, disability, age, sexual orientation, or socioeconomic status, whether that conduct is directed to other counsel, court personnel, witnesses, parties, judges, judicial officers, or any persons involved in the legal process*

The use of generative AI in the practice of law, assistance in making judicial decisions, in the administration of a firm as a business, and otherwise risks adopting biases built into the models as a result of their training data. While it is unclear if Rule 8.4(g) would be used to discipline a lawyer for subtle bias in work product or mere disparate impact, preventing model bias is also important to provide accuracy in results.

#### **Recommended Practices:**

- Ensure the prompts or other information provided to generative AI tools contain sufficient relevant data points to minimize the risk of a model latching on to prejudices or stereotypes in producing output.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Monitor AI output for evidence of bias polluting the result.

### 2. The Unauthorized Practice of Law

Generative AI can mimic some of what a lawyer normally does in answering legal questions, drafting legal documents, and providing advice. The emergence of consumer-facing generative AI tools that can answer legal questions, generate legal documents, and offer context-specific guidance has created challenges for the traditional framework governing UPL. This part addresses the obligations of licensed lawyers who use AI in ways that may create or facilitate UPL.

With some limited exceptions, only lawyers licensed or authorized by the Colorado Supreme Court can practice law in Colorado, and the Colorado Supreme Court regulates the practice of law by those who are not authorized. The Office of Attorney Regulation Counsel (OARC) investigates and prosecutes complaints alleging the unauthorized practice of law (UPL).

#### ***Rule 232.2 Governing Unauthorized Practice of Law Proceedings***

*“Practice of law” includes the following:*

- (1) Protecting, defending, or enforcing the legal rights or duties of another person;*
- (2) Representing another person before any tribunal or, on behalf of another person, drafting pleadings or other papers for any proceeding before any tribunal;*
- (3) Counseling, advising, or assisting another person in connection with that person’s legal rights or duties;*
- (4) Exercising legal judgment in preparing legal documents for another person; and*
- (5) Any other activity the supreme court determines to constitute the practice of law.*

...

*The unauthorized practice of law by a nonlawyer includes the following:*

...

- (9) Owning or controlling a website, application, software, bot, or other technology that interactively offers or provides services involving the exercise of legal judgment[.]*

The Colorado Supreme Court defines the "practice of law" as acting "in a representative capacity in protecting, enforcing, or defending the legal rights and duties of another and in counseling, advising and assisting [another] in connection with these rights and duties." Colorado Rules of Civil Procedure (C.R.C.P.) 228 to 240 govern the unauthorized practice of law and set forth the remedies available to protect against those engaging in UPL. Section (c)(9) specifically notes that the use of technology by a nonlawyer to provide legal advice is expressly defined as a prohibited unauthorized practice of law.

There are some exceptions carved out by caselaw in Colorado for limited situations such as licensed professionals acting within specific frameworks. In general, though, the UPL rules prohibit non-lawyers from practicing law even if they do it using generative AI technology.

***Rule 5.5. Unauthorized Practice of Law; Multijurisdictional Practice of Law***

*(a) A lawyer shall not:*

*(1) practice law in this jurisdiction without a license to practice law issued by the Colorado Supreme Court unless specifically authorized by C.R.C.P. 204, et seq. or federal or tribal law;*

*(2) practice law in a jurisdiction where doing so violates the regulations of the legal profession in that jurisdiction;*

*(3) assist a person who is not authorized to practice law pursuant to subpart (a) of this Rule in the performance of any activity that constitutes the unauthorized practice of law;*

A licensed lawyer, by contrast, can practice law. But, even here, the lawyer must be cautious. Generative AI deployed by a lawyer may allow the lawyer to reach more clients or prospective clients, increasing a risk that the lawyer might be exceeding what is permitted.

Under Rule 5.5, licensed lawyers who develop, deploy, market, or partner with AI legal service providers must ensure that the service does not cross into unauthorized practice. Similarly, a lawyer may not normally delegate the task of providing legal advice to others or assist others in the unauthorized practice of law.

Under Rule 5.5(a), a lawyer who assists a company in providing AI-based legal services without adequate lawyer supervision may share responsibility for any resulting UPL.

***Safe Harbor Provision***

A central principle in UPL analysis is that nonlawyers may perform tasks that would otherwise constitute the practice of law if those tasks are performed under the active supervision of a licensed lawyer who bears responsibility for the work. This principle likely applies to AI tools: an AI-generated legal document or piece of advice that is reviewed, verified, and adopted by a supervising licensed lawyer generally does not constitute UPL, even if the underlying generation was performed by the AI.

## Colorado LTAC AI Guidance Document — Synthesized Draft

OARC has also adopted a non-prosecution policy stating it will generally refrain from prosecuting UPL under certain conditions where a licensed lawyer is supervising the work of the nonlawyers. This policy likely applies to generative AI tools as any other nonlawyer use. As currently constituted, this policy only applies where a long series of conditions are true. Those conditions and the policy itself may change over time, and so anyone concerned with this policy should seek out the most current version in force.

As of the date of this guidance, the policy generally does not apply to juvenile or adult imprisonment, the safety of a minor, or the termination of parental or custodial rights of a minor or where there is no lawyer supervising the work or acting as compliance officer. Additionally, the policy has requirements for the disclosure and representation of the services being offered.

Additionally, the non-prosecution policy only applies if the nonlawyer is not holding themselves out as representing anyone or concealing the nature of their services, not made any statements deceptive under law or suggesting the non-lawyer's services are equal to those from a lawyer or LLP, doesn't cause harm through gross negligence or more culpable standard, provides all the services paid for, doesn't harm a consumer by reducing legal rights or remedies without warning, and has not prejudiced justice.

The policy, finally, explains that to avoid prejudicing justice, the non-lawyer must take reasonable steps to avoid output containing errors, hallucinated cases, false evidence, or other similar output generated by the nonlawyer, warn users about such errors, hallucinations, false recitations or similar output, and not knowingly allow users to abuse the legal system.

While the particulars of the non-prosecution policy may be subject to change in the future, the general guidelines and standards that would apply to using generative AI tools to provide legal advice are generally consistent with the practices explained elsewhere in this guidance document. Lawyers who build, supervise, or partner with such tools should track its current terms.

### **Recommended Practices for lawyers who supervise, build, or partner with AI legal services:**

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Ensure that any AI legal service they supervise or endorse includes meaningful lawyer review for matters in which legal judgment is applied to specific client facts; and
- Not lend their license or professional identity to AI services in a manner that creates a false impression of lawyer supervision when none exists.
- Where relying on OARC's non-prosecution policy rather than individual lawyer review, confirm that the service meets the policy's requirements, including its subject-matter exclusions, keep abreast of any change to the policy or related rules, and have a plan to come into compliance with the UPL rules should the policy sunset or change.
- Be aware that multijurisdictional issues arise when AI tools deployed in one state provide advice affecting users in other states where the supervising lawyer is not licensed. The non-prosecution policy is limited in jurisdiction to the state of Colorado.

### 3. The Judiciary

Judges and judicial officers face a distinct set of obligations when using generative AI. The Colorado Code of Judicial Conduct governs not only the substance of judicial decisions, but the integrity of the judicial process itself. The use of AI tools by judicial officers raises issues of impartiality, *ex parte* communications, independence, and the public's perception of the judiciary. This unit discusses the canons and rules most directly implicated by judicial use of AI, and then the practical boundaries on administrative use by judges and court staff.

#### ***Canon 1. Independence, Integrity, and Impartiality***

*Canon 1. A judge shall uphold and promote the independence, integrity, and impartiality of the judiciary, and shall avoid impropriety and the appearance of impropriety.*

AI tools trained on legal texts may reflect biases that could affect the perceived or actual impartiality of AI-assisted judicial work product. Judges who use generative AI tools in drafting opinions, conducting legal research, or analyzing facts must be alert to the risk that AI outputs may reflect embedded biases in the tool's training data. These biases are typically not visible to the user and may be difficult to detect.

#### **Recommended Practices:**

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Exercise independent judgment in reviewing and evaluating AI-generated research and drafts. Do not treat them as authoritative.
  - Be aware that AI tools trained predominantly on materials from particular jurisdictions or legal traditions may not accurately reflect Colorado law.
  - Remain alert to the possibility that AI tools may generate text that subtly reflects particular legal, political, or cultural viewpoints embedded in training data.
  - Ensure that the ultimate analysis and judgment in judicial decisions is the product of the judicial officer's independent reasoning, not AI output.

***Canon 2. Performing the Duties of Judicial Office  
Impartially, Competently, and Diligently***

*Canon 2. A judge shall perform the duties of judicial office impartially, competently, and diligently.*

Just as lawyers must understand the tools they use, so too must judges. A judicial officer who relies on AI-generated legal research without adequate understanding of the tool's limitations risks basing decisions on inaccurate information.

### **Recommended Practices:**

- Complete training programs on the use of generative AI tools in judicial work before incorporating them into chambers practice, and support judicial education programs on AI.
  - Verify AI-generated legal citations and propositions with the same rigor required of lawyers submitting court filings.
  - Ensure that law clerks and chambers staff using AI tools are adequately trained and that their AI-assisted work is subject to appropriate judicial review.

Law clerks and judicial staff are increasingly using AI tools to assist with bench memos, opinion drafts, and legal research. The judge, as the supervising judicial officer, bears ultimate responsibility for the accuracy and integrity of chambers work product.

### **Recommended Practices:**

- Establish clear chambers policies on which AI tools may be used, under what conditions, and with what verification requirements.
  - Require clerks to disclose to the judge when AI tools have been used in preparing bench memos or draft opinions.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Review AI-assisted research and drafts with the same critical attention applied to traditional work product.
- Ensure that final judicial work product reflects the judge's own analysis and reasoning, not merely edited AI output.

### **Rule 2.9. Ex Parte Communications**

*(A) A judge shall not initiate, permit, or consider ex parte communications, or consider other communications made to the judge outside the presence of the parties . . .*

When a judicial officer inputs case-specific facts, arguments, or evidence into an external AI platform, several risks arise:

- The AI platform may process and store the information in ways that could theoretically be accessible to third parties.
- The AI tool's response may be influenced by its training data in ways that amount to consideration of information not in the record.
- Using an AI tool to explore legal arguments not raised by the parties may raise ex parte concerns if the tool is trained on external advocacy materials that effectively represent particular legal positions.

### **Recommended Practices:**

- Avoid inputting case-specific party names, confidential record information, sealed materials, or substantive case facts into any external AI platform not specifically vetted and approved for use in judicial work.
- Frame queries to AI tools as general legal research questions, not as case-specific inquiries using actual case facts or identifying information.

### **Rule 2.10. Judicial Statements on Pending and Impending Cases**

*(A) A judge shall not make any public statement that might reasonably be expected to affect the outcome or impair the fairness of a matter pending or impending in any court . .*

Judges must exercise caution when using AI to draft any communications that touch on matters pending before their courts or matters of current legal controversy. AI tools may generate text that, while technically accurate, creates an appearance of prejudice or partiality. All AI-assisted public communications should be carefully reviewed for compliance with Rule 2.10 before release.

**Rule 2.11. Disqualification**

*(A) A judge shall disqualify himself or herself in any proceeding in which the judge's impartiality might reasonably be questioned . . .*

As AI companies become parties to litigation, and as cases involving AI-generated products, copyrights, and harms proliferate, judges must be attentive to potential conflicts arising from:

- Personal investments in AI companies;
- Institutional court contracts with AI vendors;
- Prior professional relationships with AI companies or their counsel; and
- Situations in which the judge's use of a particular AI tool creates an appearance of endorsement or partiality.

**Administrative Use and the Adjudicative Core**

The judiciary occupies a critical role in maintaining the integrity of the legal system while adapting to technological advances. Judges, magistrates, and court administrators must balance the administrative efficiencies offered by AI with the strict requirement to maintain judicial independence, impartiality, and the security of the court's records.

AI may be used as a tool to assist with the judge's time and administrative burden, but it must never replace the judge's independent reasoning or adjudicative process. The core functions of the bench, such as weighing conflicting evidence, assessing the credibility of witnesses, determining the equitable distribution of assets, or interpreting the application of law to novel facts, are exclusively human endeavors that must never be delegated to an algorithm.

Drafting deserves a specific word here. Though AI tools are competent drafters, the act of writing an order can itself sharpen judicial decision-making: working through the words is often how a judge works through the decision. Judges who use AI to assist with drafting orders should thoughtfully develop processes that ensure they still move through the decision-making steps methodically and thoroughly, rather than ratifying a draft that arrived already written.

Judicial officers may safely and effectively use approved, secure AI tools for administrative tasks. Examples include generating timelines from voluminous exhibits, summarizing lengthy hearing transcripts, or proofreading draft opinions for grammatical

## Colorado LTAC AI Guidance Document — Synthesized Draft

consistency. However, even in these administrative contexts, the judge or the judge's law clerks must verify AI output against the official record.

Courts must also be prepared to address instances where litigants, both represented and self-represented, submit filings containing AI-generated content. Courts may consider adopting local standing orders that remind practitioners of their verification duties and require human review of all submissions. If a court suspects a citation is hallucinated, it is within the court's discretion to order the filing party to produce a copy of the cited authority.

For court administrators and clerks, data security is the top concern. Court staff must strictly follow internal protocols, ensuring that non-public, sealed, juvenile, or otherwise confidential records are never processed through unauthorized public AI models. When interacting with self-represented litigants, court staff may use AI tools internally to locate correct forms or general procedural guidelines, but they must be careful not to pass off AI-generated responses as personalized legal advice, maintaining the strict boundary between legal information and legal practice.

#### **4. Licensed Legal Paraprofessionals.**

Licensed Legal Paraprofessionals (LLPs) are a new licensed legal profession in Colorado. Under the C.R.C.P. 207 series, LLPs may provide certain legal services in certain family-law matters without a law license.[28] LLPs are not lawyers, but they are not unregulated either. They are subject to the Colorado Licensed Legal Paraprofessional Rules of Professional Conduct (the LLP RPCs), a separate rule set modeled on the lawyer rules.

Because the rules parallel each other, nearly everything said about lawyers in Sections 1 and 2 of this Part applies to LLPs with equal force. This is true even though the LLP RPC do not include the same AI-specific comment to Rule 1 that is also found in the Colorado Rules of Profession Conduct for lawyers. The duty of competence in LLP RPC 1.1 reads the same way the lawyer rule does, generative AI is "relevant technology" for an LLP's practice just as it is for a lawyer's, and an LLP remains accountable for work product regardless of the technology that produced it.

**Colo. LLP RPC 1.1. Competence**

An LLP shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary to:

- (a) perform the contracted services; and
- (b) determine when the matter should be referred to a lawyer.

An LLP who uses generative AI without understanding its capabilities and limits risks providing incompetent representation, the duty does not require avoiding AI, and the duty is to use the tools with knowledge, skill, and appropriate verification. But, there are additional concerns when tools are used by an LLP.

LLP license is examined and defined around family law, a very specific kind of law, and AI tools are not. A general-purpose model will answer a family-law question with the same confidence it answers anything else, including with out-of-state law, superseded statutes, or procedures from the wrong court. Many of the tools an LLP is likely to encounter are the AI features bundled into practice-management and office software. Competence includes knowing which of those features are turned on, what they do with client data, and whether their output is being checked, rather than treating them as part of the furniture.

**Recommended Practices:**

- Obtain adequate training before using generative AI in client matters, and treat AI output as a first draft requiring the LLP's own substantive review, editing, and judgment.
- Test tools on family-law problems where you already know the answer — the forms, deadlines, and statutes you work with daily — using the evaluation methods in Part 4.
- Inventory the AI features embedded in your practice-management and office software, and configure them deliberately.

**C.R.C.P. 207.1 and Colo. LLP RPC 5.5**

C.R.C.P. 207.1

*(2)(a) An LLP may represent a client to perform tasks and services identified under section (2)(g) of this rule in a legal separation, declaration of invalidity of marriage, or dissolution of a marriage, or and dissolution of a civil union, whether pre-decree or post-decree [and] to perform [certain] tasks and services... in an initial allocation or a modification*

Colo. LLP RPC 5.5

*An LLP shall not...*

*(1.5) practice law beyond the authorization set forth by the Colorado Supreme Court...*

*(3) assist a person who is not authorized to practice law in the performance of any activity that constitutes the unauthorized practice of law*

Generative AI does not expand the LLP license. The scope of what an LLP may do is fixed by C.R.C.P. 207.1, not by what an LLP's tools can produce. A frontier model can draft a civil complaint, a criminal motion, or an immigration filing as fluently as it drafts a parenting plan. That capability is a trap. An LLP who uses AI to take on work outside the authorized scope is practicing law without authority, no matter how good the output is, and exceeding scope can violate the LLP RPCs directly.[31]

Family-law matters regularly raise adjacent issues like tax consequences of a property division, immigration effects of a decree, criminal exposure from a protection-order violation and a model asked about a family-law problem will happily analyze the adjacent issue as well, because the model has no concept of a license boundary. The LLP must supply the boundary the tool lacks: recognize when an answer has crossed into territory that requires a lawyer, stop, and refer.

The same logic applies in the courtroom. The 207.1 limits on what an LLP may do before the court are not loosened because an AI tool prepared the material.

**Recommended Practices:**

- Before relying on any AI output, ask whether the task — not just the document — is within the scope of C.R.C.P. 207.1. If it is not, the polish of the output is irrelevant.
- Treat AI analysis of issues adjacent to a family-law matter (tax, immigration, criminal, bankruptcy) as a signal to consider referral to a lawyer, not as a license to advise on them.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- When marketing or describing your services (see LLP RPC 7.1 below), never suggest that AI tools let you handle matters beyond your authorized scope.

- 

### **Colo. LLP RPC 1.6. Confidentiality of Information**

*An LLP shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized in order to carry out the representation, or the disclosure is permitted[.]*

The full Rule 1.6 analysis in Section 1 for lawyers including vetting data practices before any upload, avoiding public-facing tools for client-identifying information, the unsettled privilege question applies to LLPs without modification.

In addition, LLPs should remember that family-law files are dense with exactly the information this guidance warns about most: domestic-violence allegations, children's information, financial statements, addresses, and immigration status. The “strip anything identifying” discipline and the preference for written no-training commitments described in Part 4 are not optional extras in this practice area; they are the baseline.

The LLP RPCs permit LLPs to practice in firms without lawyers.[32] An LLP practicing independently may be more likely to have no firm governance layer: no IT department vetting tools, no committee approving them, and no supervising lawyer to catch a misconfigured setting. Like the solo lawyers discussed in Part 4, independent LLPs carry the entire evaluation and confidentiality burden personally, and the vendor-diligence checklist in Part 4 is written for exactly this situation.

### **Recommended Practices:**

- Apply the Rule 1.6 checklist from Section 1 before entering any client information into any AI tool, and prefer tools with written no-training commitments for family-law material.
  - If practicing without lawyers in the firm, adopt a written personal AI-use policy: which tools are approved, what data may enter them, and what verification happens before anything is finalized.
  - Be especially cautious with matters involving domestic violence, children, or immigration status. Consider whether such matters should be handled with no third-party AI exposure at all.

**Colo. LLP RPC 3.1 and 3.3. Meritorious Claims and Candor Toward the Tribunal**

Colo LLP RPC 3.1

*An LLP shall not bring or defend a proceeding, or assert or controvert an issue therein, unless there is a basis in law and fact for doing so that is not frivolous, which includes a good faith argument for an extension, modification or reversal of existing law*

Colo LLP RPC 3.3

*An LLP shall not knowingly ... make a false statement of material fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the LLP.*

An LLP who signs and files a document answers for it exactly as a lawyer does. Fabricated citations, misrepresented holdings, and invented authority in an LLP's filing violate the LLP RPCs, expose the LLP to discipline up to loss of license, and harm the client. The verification duty is identical as with lawyers. An LLP must confirm that every cited case, statute, rule, and JDF form reference exists, says what the filing claims, and remains current. Nothing about the limited scope of the license limits this duty.

**Recommended Practices:**

- Verify every citation and form reference in a trusted source before filing — the model is never the source.
- Where a court's standing order addresses AI use or requires certification, comply with it precisely; the trackers in the Annex collect current orders.

**Colo LLP RPC 5.3 Responsibilities Regarding Other Personnel in Firms Without Lawyers**

- (a) *an LLP who individually or together with other LLPs possesses comparable managerial authority in a firm without lawyers shall make reasonable efforts to ensure that the firm has in effect measures giving reasonable assurance that the person's conduct is compatible with the professional obligations of the LLP;*
- (b) *an LLP having direct supervisory authority over other personnel shall make reasonable efforts*

An LLP who supervises other LLPs, paralegals, or staff bears the same supervisory duties for their AI use that Sections 1 and 2 describe for lawyers. The LLP should have clear policies on permitted tools and data, training, and verification before anything goes out. A subordinate's undisclosed AI use is the supervising LLP's problem.

## Colorado LTAC AI Guidance Document — Synthesized Draft

For unlicensed paralegals, a reminder: an unlicensed paralegal has no authority to practice law and must be directly supervised. A paralegal who uses a generative AI tool to exercise legal judgment for a client has not found a loophole. Under C.R.C.P. 232.2, interactive technology offering services involving legal judgment is itself defined as the unauthorized practice of law, and the supervising lawyer or LLP shares responsibility for letting it happen.

### Recommended Practices:

- Firms employing LLPs: extend the firm AI policy, training, and verification requirements explicitly to LLPs, and to the boundary between in-scope and out-of-scope work.
- Supervising LLPs: ask whether AI was used when reviewing subordinates' work, and build a culture where staff can report suspected AI errors without fear.

#### ***Colo. LLP RPC 7.1 Communications Concerning an LLP's Services***

*An LLP shall not make a false or misleading communication about the LLP or the LLP's services. A communication is false or misleading if it contains a material misrepresentation of fact or law, or omits a fact necessary to make the statement considered as a whole not materially misleading*

The Section 1 discussion of Rule 7.1 applies, with one addition specific to the LLP market. LLPs compete in part on affordability, and they may be tempted to advertise how AI efficiency contributes to that affordability. But an LLP's communications must not overstate what AI tools deliver, must not suggest that AI lets the LLP provide services equivalent to a lawyer's beyond the authorized scope, and must keep any claims about AI-driven savings or capabilities current as the tools change.

### Recommended Practices:

- Do not over-promise AI benefits, cost savings, or reliability, and review marketing materials whenever your tools change.
- Never describe AI capabilities in a way that blurs the line between LLP services and a lawyer's.

**Colo. LLP RPC 1.5 Fees**

*An LLP shall not make an agreement for, charge, or collect an unreasonable fee or an unreasonable amount for expenses.*

The fee analyses in Section 1 apply in full: bill the time actually spent, not the time the tool saved; remember that prompt-crafting and verification are real billable work; remember subscriptions may be unbillable overhead. The fee points carry extra weight for LLPs, whose value to clients rests on limited-scope, lower-cost service. AI efficiency gains may be proportionally larger in a practice model where flat-fee arrangements are common, and honest pricing of AI-assisted work is close to the core of what the LLP program exists to deliver.

**Colo. LLP RPC 8.4(g)**

*It is professional misconduct for an LLP to... engage in conduct, in the representation of a client, that exhibits or is intended to appeal to or engender bias against a person on account of that person's race, gender, religion, national origin, disability, age, sexual orientation, or socioeconomic status, whether that conduct is directed to other LLPs, counsel, court personnel, witnesses, parties, judges, judicial officers, or any persons involved in the legal process.*

The bias analysis in Section 1, finally, also applies in full: AI output should be monitored for bias, with prompts supplying enough context to keep the model from pattern-matching to stereotypes.

## **Part 6 : Self-Represented Litigants**

It is unlikely that an AI tool will be as effective as hiring a good lawyer or LLP. Nevertheless, particularly if doing this is cost prohibitive, AI tools can help you understand your rights, fill out legal forms, and prepare for court better than you could on your own. But not all AI tools are safe to use for legal problems, and even the best ones have real limits. This guide explains what to look for, what to watch out for, and what your rights are when you use AI for legal help. Many of the ideas here are expanded upon in Part 4 – Resources earlier in this Guidance document.

### **1. When AI Helps and When to Get a Person**

Since AI tools do have problems, you might be tempted to avoid them entirely. That may not always be the right choice. For many people, the realistic alternative to an AI tool is no help at all. So the goal is not to swear off these tools. The goal is to use them well and to know when your situation calls for a person instead.

Using AI tools as a self-represented litigant represents a tradeoff. Used carelessly, AI can hurt your case. Used well, it can help. The following factors move the needle, and weighing them is how you decide how far to trust a tool and when to stop and get human help.

- **Is a licensed lawyer reviewing the work?** Output that a licensed attorney has reviewed, corrected, and stands behind is far safer than output no lawyer has ever looked at. A tool built around real lawyer review is closer to getting actual legal help than to using a chatbot alone.
- **How high are the stakes, and can a mistake be undone?** Losing your home, your children, your immigration status, or your freedom is not something you can easily reverse. The higher the stakes and the harder a mistake is to fix, the more you should lean toward getting a person involved.
- **Can you realistically check the work?** Some answers you can verify yourself, such as confirming a form, a deadline, or whether a cited case exists. Others require legal judgment you may not have. The less able you are to catch an error, the more a serious matter calls for human help.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- **What area of law is this?** Some areas are high-risk by their nature. Colorado's own rules treat juvenile cases, criminal charges, and cases about ending parental rights as areas where caution matters most. In those areas, get human help.

None of these factors decides the question by itself. Read together, they give you a way to judge, for your own situation, whether an AI tool is enough or whether it is time to bring in a person.

### 2. Two Very Different Types of AI Tools

If you are thinking about using an AI tool, one of the first things you need to know is that there are two very different kinds of AI tools, and they are not equally safe for legal problems.

#### **2(a). General-Purpose AI Platforms (Use With Caution)**

Tools such as ChatGPT, Google Gemini, and Claude were not designed for legal help. They can sound confident and helpful, but they have serious problems for legal use:

- They do not know Colorado law specifically.
- They can give you outdated information without telling you it is outdated.
- They can make up court cases, statutes, and legal rules that do not exist. This is called "hallucination." The answer looks real but is not.
- Even if they give you real law or cases, they might not understand how to apply them correctly.
- They have no obligation to tell you when they do not know something.
- You might be waiving certain confidentiality and privilege protections by sharing information with the AI developer.
- Your conversation and your personal information may be stored and used to train the AI in the future.

If you use one of these tools for a legal problem, know that it comes with some serious risks. If you can, verify everything it tells you before you act on it. Better yet, use a purpose-built legal tool instead.

#### **2(b). Purpose-Built Legal Assistance Tools (What to Look For)**

A purpose-built legal tool is designed specifically for legal help. A good one is built around Colorado law, reviewed by licensed attorneys, tested with real users, and designed

## Colorado LTAC AI Guidance Document — Synthesized Draft

to tell you when it cannot help. These kind of tools are probably more likely to be grounded in real law and have higher standards for confidentiality.

Part 4 of this Guidance document provides detailed lists of questions to ask and resources to consult concerning how to choose a good tool to help you. If you are seriously considering using an AI tool to assist with a legal matter, consider all of the resources explained in that Part to help minimize bad results.

### ***2(c). Questions to Ask Before You Use an AI Legal Tool***

Generally speaking only a lawyer or licensed legal paraprofessional (LLP) can actually give legal advice. There are exceptions and those exceptions may be changing over time as technology changes. Before you rely on any AI tool for a legal problem, see if the tool or its provider answers these questions. A trustworthy tool should be able to answer all of them clearly.

*Is a licensed attorney involved?*

A responsible tool should have a licensed attorney supervising the legal content, or a lawyer serving as a legal compliance officer for the organization that built it. If there is no attorney involvement, be very careful. If a provider claims that it is giving you legal advice or is as good an attorney or licensed legal professional, you probably should avoid that tool.

*Does it tell you clearly that it is not a lawyer?*

Before you start using the tool, it should tell you in plain language that you are talking to an AI system, not a lawyer. It should also tell you that using it does not create an attorney-client relationship. This warning should be simple enough that a fifth-grader could understand it, and it should appear before you type anything — not buried in fine print or a terms-of-service agreement you have to hunt for.[27] If this disclosure is hidden or missing entirely, that is a red flag.

*Does it tell you what area of law it covers or what specific legal tasks it can accomplish?*

A good tool is specific about what it handles and what it does not. It should tell you which legal issues fall within its scope and what it will do if your question is outside that scope. It should also tell you when its legal information was last updated.

## Colorado LTAC AI Guidance Document — Synthesized Draft

*Does it show you where its information comes from?*

Every legal claim the tool makes should point to a real source: a Colorado statute, a court rule, or a case you can look up. If the tool just gives you answers with no sources, you have no way to verify what it says. Remember, a judge will not care what the AI tool thinks or believes. A judge will care about the actual law. An AI tool that helps you locate the specific statute, regulation, or case that controls your case is going to be more effective than one that simply produces an argument that sounds persuasive to a non-lawyer.

*What happens to your information?*

The tool should tell you, in plain language, what it does with the information you share. Specifically, you should be told:

- What information the tool collects from you.
- Whether it is stored after your session ends.
- Who can access it.
- Whether it will be shared with anyone.
- Whether it will be used to train or improve the AI.

A responsible tool will tell you all of this clearly and ask for your agreement before you start. If that disclosure is not there, or if it is written in dense legal language you cannot understand, proceed with caution. This matters especially if your situation involves immigration status, domestic violence, or other sensitive facts. In those cases, think carefully before sharing personal details with any AI tool.

An important note: the law on whether your communications with an AI legal tool are protected the same way your conversations with a lawyer are protected is not yet settled in Colorado. Until that question is resolved, you should assume that what you share with an AI tool is not automatically confidential the way a conversation with your attorney would be. The tool should disclose exactly what protections it does and does not provide.

*Does it tell you when it cannot help?*

A trustworthy tool should recognize its limits and clearly tell you when your situation is outside its range. It should also point you to real help, such as a legal aid office or a court self-help center, when you need more than the tool can provide.

*Is it endorsed by a court or legal aid organization?*

Tools that carry the visible backing of courts, legal aid organizations, bar associations, or the Access to Justice Commission have been evaluated to some degree. That is a meaningful signal. Look for those endorsements, and be cautious about tools with none.

If you have any concerns about whether the vendor is allowed to help you in this way, you can contact the Colorado Attorney Regulation Counsel for more information about what kind of tools are acceptable versus not acceptable.

**2(d). Use the Tool as a Way to Learn, not a Robo-Attorney**

If you just blindly trust an AI tool, you are not getting the most benefit out of it. If time permits, using generative AI to learn and understand the law is likely to far more powerful. An AI tool can help explain a legal document or a law in plain language. If you do not understand the explanation, you can ask the tool follow up questions, for analogies, or for more information and background so you can better understand the law.

And, if you are not sure what questions to ask, generative AI has the unique ability to help you formulate them. If you are not sure where to start, you can describe the legal situation to a model and then ask it what questions you should be asking to understand the situation. For example, you might upload a copy of a complaint filed against you into a (properly vetted) AI model and ask it, “What questions should I be asking and what information should I be collecting to understand this complaint and what to do next?” You can ask follow up questions. You can back up and start again. A model is infinitely patient and will take as long as you need to help you understand something.

You can be more creative than just analyzing legalese, too. If you are not a lawyer, there are many conventions and expectations that lawyers and judges might know, but you may not. You can simply ask an AI tool what those conventions are and what to expect. For example, if you are due to appear in court for your first appearance in a divorce, you could ask the model, “In addition to the legal requirements, are there any other expectations or conventions that I should know about?” Similarly, if you are responding to a demand letter from an attorney, you might ask the model, “If I want to end the case by paying a certain

amount of money, how should I frame the offer and what other things should I be looking out for?”

At the end of the day, if you used an AI tool to blindly fumble through a legal proceeding, you gained nothing from its use beyond that case. But, if you use the AI tool to learn and understand the system better, you have improved yourself along the way.

### **3. Risks You Should Know About**

Even good AI tools carry risks. Here is what to watch out for.

#### **3(a). *AI can make things up***

AI systems can generate case names, statute citations, and legal rules that simply do not exist. They do this confidently, in professional-sounding language, with no warning. If you file something in court that relies on a made-up case, you can lose your case and potentially face other consequences. Always verify citations before using them.

#### **3(b). *Confident does not mean correct***

AI tools produce polished, authoritative-sounding answers even when they are wrong. That confident tone can make you feel like you do not need to double-check or get a second opinion. That is one of the most dangerous features of AI in legal situations. If an answer feels too clear and easy, that is a reason to verify it, not a reason to stop looking.

#### **3(c). *The tool may not know what it does not know***

Unlike a lawyer, who knows when a question is beyond their area of expertise, an AI tool may not recognize the limits of its own knowledge. It may miss a defense you have, apply the wrong state's law, or rely on a rule that changed last year. You cannot always tell when this is happening. This one reason why it is so important to use an AI tool to find and understand the law, not just to give you conclusions you blindly trust.

#### **3(d). *The tool may not ask the right questions***

A good AI tool should surface rights and defenses you may not know to ask about. But some tools are purely reactive and only answer what you ask. If you do not know the

## Colorado LTAC AI Guidance Document — Synthesized Draft

right questions, you may not get the most important information. For example, in an eviction case, you might have a defense based on your landlord's failure to maintain the property, but if the tool does not prompt you about that, you might not raise it.

### **3(e). *The tool may push you toward one choice***

The decisions in your case belong to you. A good tool lays out your options and explains the tradeoffs of each. Be wary of a tool that steers you toward a single course of action — for example, nudging you to settle or to drop a claim — without explaining the alternatives. Even if the suggested path is reasonable, a tool that decides for you has taken away the choice that is yours to make. If a tool keeps recommending one path, ask it directly what your other options are, and what you would give up by taking each one.

### **3(f). *Procedurally correct is not the same as strategically sound***

An AI tool might help you file something correctly that is still the wrong strategic choice for your situation. Knowing how to file a document is different from knowing whether filing it is the right move. AI tools do not replace the judgment of a lawyer who knows your full situation.

### **3(g). *Your information may not be private***

When you share personal information with an AI tool, whether it is your address, your immigration status, the details of a domestic violence situation, or facts about a court case, you may be putting that information into a system whose data practices you do not fully understand. Section 5 explains your privacy rights.

### **3(h). *You are least able to catch the mistakes you can least afford***

The people who most need AI legal help are also the people least equipped to spot AI errors. Legal training is what lets a lawyer know when an answer is wrong. If you had that training, you would not need the tool. This is not a reason to avoid AI tools. It is a reason to understand their limits and to get human help when the stakes are high.

## **4. What to Do When a Document Will Be Filed in Court**

## Colorado LTAC AI Guidance Document — Synthesized Draft

If you are using an AI tool to help you prepare something that will be filed in court, take these extra steps. If you do not, you could be punished by a Court by losing a motion, being forced to pay a large fine, or otherwise.

- Read the document carefully before you file it. Make sure it accurately describes your situation.
  - Verify every legal citation. Look up any case or statute the document refers to and confirm it exists and says what the document claims.
  - Check your deadline. Filing after your deadline may mean losing your right to respond or raise a defense.
  - If anything is unclear or the stakes are high, contact a legal aid organization or a court self-help center before you file.

### 5. Where to Get Help Beyond AI Tools

Colorado has a number of free or low-cost resources for people who need legal help.

- LawHelp Colorado: Free legal information and referrals for low-income Coloradans.
- Court self-help staff: Some courthouses have self-help staff, known in some Colorado courthouses as Sherlock staff, who can help you navigate the court process and find resources.
  - Legal aid organizations: Organizations such as Community Economic Defense Project and Colorado Legal Services provide free legal representation to people who qualify.
  - Law school clinics: Several Colorado law schools run clinics that provide free legal assistance to community members.
  - Public libraries: Many libraries can help you access legal resources and may have staff who can connect you to legal help.

A well-designed AI legal tool will refer you to these resources when your situation is beyond what it can reliably handle. If a tool does not do that, that is a sign that it may not have your best interests in mind.

Colorado courts are working out how to serve people who represent themselves in a world where AI tools exist. Here is what that can look like for you.

## Colorado LTAC AI Guidance Document — Synthesized Draft

- Courts and legal aid organizations evaluate some AI tools and may identify tools that have been found accurate and safe. When a court or legal aid organization visibly endorses a tool, that endorsement is a meaningful signal you can rely on.
  - Court self-help centers and self-help staff can help you understand whether documents prepared with an AI tool meet the court's filing requirements, including its electronic filing system.
  - Court staff and judges are receiving training about AI legal tools, including the responsible ones. If you used a responsible tool carefully and verified your work, you have done what the rules ask of you.

### Quick Reference: The Bottom Line

- **Not all AI tools are safe for legal problems. Use purpose-built legal tools over general-purpose AI.**
- **Any tool you use should tell you it is not a lawyer, in words anyone can understand, before you start.**
- **Responsible tools should tell you that a lawyer reviewed the legal content and when it was last updated.**
- **Never file anything in court without carefully reading it and verifying any legal citations.**
- **Know what happens to your personal information before you share it.**
- **The decisions in your case are yours. A good tool shows you options; it does not choose for you.**
- **The higher the stakes, the more you should get a real person involved.**