Legal Technology Advisory Committee
Subcommittee on AI and Self–represented parties
2.16.2026

Outline of Issues

Note: Many of these sub-topics overlap and will implicate similar issues, guidance and discussion. Feel free to consolidate and/or add sub-topics as appropriate.

I. The risks AI poses to self-represented parties and recommendations for minimizing these risks.

Examination of the normative questions of which principles (justice, integrity, equality agency, transparency, privacy, and fairness) should be paramount and prioritized when we provide guidance to self-represented parties.

- We should be able to articulate an explicit ordering of principles that reflects the conditions of self-represented party interactions. Here is one prioritization framework:
    - Justice as the overarching goal/ultimate objective.
        - Just want to make sure we're focusing on the access to justice gap rather than some amorphous sense of justice. Regarding the former, AI tools have the potential to begin correcting the disparity. But to do so, the guidance governing AI tools must be designed with that corrective purpose explicitly and unambiguously in mind, or the tools will simply reproduce existing inequalities in a more technologically sophisticated form.
    - Integrity as the operational prerequisite.
        - Judicial integrity is an important value, but it is not well served by restricting access to AI tools. When tools are accurate, well-designed, and honestly disclosed, they do not threaten the integrity of the judicial system. A guidance document that treats integrity as the overriding concern will tend to produce restrictions that protect courts at the expense of the litigants those courts exist to serve. The more accurate framing is that integrity and access point in the same direction when the tools work as they should.
    - Informed agency as the relational commitment to users.
        - Autonomy deserves particular emphasis as foundational. The goal is not to optimize outcomes but to empower people to participate meaningfully in proceedings that affect their lives. A tool designed to optimize outcomes might steer a litigant toward settlement even when that litigant,

fully informed, would choose to fight. That distinction has direct design implications.

- Substantive equality as the standard for evaluating systemic performance (benchmark).
    - The equality principle imposes real demands on how we design and evaluate AI guidance. A set of verification protocols or access requirements that only a tech-savvy litigant can navigate easily may be a nearly insurmountable barrier for someone with limited digital literacy, limited English, or limited time. The population most likely to need AI legal assistance is also the population least likely to be able to satisfy demanding access conditions, and guidance designed with the median or sophisticated user in mind will systematically fail the users who need it most. Guidance that is formally neutral but structurally exclusionary violates the equality principle just as surely as guidance that explicitly disadvantages particular groups. If we are not asking at every step whether a proposed requirement would work for the least advantaged user, we are not actually applying an equality standard.
- Meaningful transparency as the informational floor.
    - I think transparency serves the other values on this list rather than standing as an independent goal. Its function is to enable genuine autonomy by ensuring that users understand what the AI tools can and cannot do, and to enable accountability by making it possible to identify and correct failures. A transparency requirement that produces dense terms-of-service disclosures satisfies the letter of the principle while likely defeating its purpose. What self-represented litigants need is plain-language communication of real limitations at the moment those limitations matter, not boilerplate at session initiation. It is also worth noting that overpromising the reliability of court-endorsed AI tools is not a neutral act. When those tools inevitably produce errors, unearned confidence in their accuracy will have made the harm worse.
- Privacy as a near-absolute constraint in high-risk contexts.
    - Definitely (closer to) an absolute constraint than a preference to be balanced against other goals. Disclosure of information in immigration, domestic violence, and housing proceedings could mean deportation, physical danger, etc.
- Fairness as procedural consistency evaluated against justice outcomes.
- Respect**.** Self-represented litigants are not problems to be efficiently processed. They are people with legitimate claims on a system that belongs to them as much as to anyone else. The design of AI tools should reflect that.
- Capabilities. The relevant question is not whether self-represented litigants have the formal right to use AI tools but whether they have the genuine capacity to do so effectively. Formal access without real capability is not access in any meaningful sense.

- This ordering is not necessarily permanent or universal and should be revisited as the technology, the legal landscape, and our understanding of user impact evolve.
- Any normative framework must be operationalized, not aspirational. Each principle should map to specific, auditable design requirements.

- Hallucination and procedural error: AI systems can confidently generate incorrect filings, cite nonexistent case law, or misstate jurisdictional rules, leaving litigants worse off than if they had filed nothing at all.
- False confidence: Polished, authoritative-sounding AI outputs may cause self-represented parties to forgo human help they urgently need, with no professional backstop to catch errors.
- Systematic bias against vulnerable populations: Tools trained on majority-population legal patterns may not accurately serve communities who may have language barriers, immigration issues, or non-standard housing arrangements.
- Opacity of failure: Unlike a lawyer who knows when a question exceeds their competence, users cannot assess what an AI system doesn't know or got wrong.
- Over-reliance and automation bias: self-represented parties may defer to AI outputs even when facts are incomplete, the tool is uncertain, or the answer conflicts with court instructions
- Data exposure risk: Self-represented parties may unknowingly disclose sensitive facts, documents, or admissions into systems whose retention and use they do not understand.
- Strategic error risk: AI might prompt an unsophisticated litigant to pursue a procedurally accurate path that’s strategically unsound.
- Autonomy illusion(?): AI tools ostensibly empower litigants but may actually supplant independent judgment through confident recommendations that users feel ill-equipped to evaluate/contest.
- Verification/Correction gap: Those least able to detect AI mistakes are also likely least able to remedy them, which compounds a disadvantage for the users already most at risk.

Recommendations:

- Build tools with constrained scope rather than general-purpose legal advice — narrow, well-tested functionality is safer than broad capability.
- Implement robust uncertainty disclosure so users understand the limits of what the system knows and when to seek human help.
- Design pathways that route users to human assistance when queries fall outside system competence.
- Test rigorously against actual user populations, not proxy datasets — accuracy for the average user does not equal safety for the most vulnerable.
- Prioritize privacy-first architecture (local or self-hosted models where feasible) given the sensitivity of legal data for vulnerable populations.

- Require task-specific warnings, not generic disclaimers
- Require fact sufficiency checks before the system answers
- Present options and explain tradeoffs rather than steering users toward a single recommended course of action.

II. Guidelines for the development of guardrails for the development of (1) bespoke AI systems for self-represented parties and (2) general-purpose legal tools (e.g., legal research assistance).

I. Bespoke AI Systems for Self-Represented Parties

Scope & Competence Boundaries

- Define and enforce a hard scope boundary: the system should be designed to answer a defined set of legal questions within a defined jurisdiction, and should affirmatively decline queries outside that scope rather than attempting a general answer.
- Build in jurisdictional awareness so the system never applies the law of the wrong state, county, or court without explicit disclosure.
- Calibrate guardrails to stakes: tools used in family law, immigration, or other high-consequence matters should be subject to stricter accuracy and oversight requirements

Transparency & Disclosure

- Require plain-language uncertainty flags whenever the system is operating near the edge of its training data, in areas of unsettled law, or in response to a query that doesn't fit its core use case.
- Mandate source citation for all legal claims, with citations linking to actual, verifiable primary sources.
- Display prominent, non-dismissible disclosures at session initiation identifying what the tool can and cannot do.

User-Centered Design Safeguards

- Conduct plain-language readability testing with actual members of the intended user population before deployment.
- Build in active referral pathways. When a user's situation exceeds system competence, the tool must affirmatively route to legal aid, court self-help centers, or other human resources rather than simply disclaiming liability.
- Design for low-literacy and multilingual users from the outset, not as an accessibility add-on.
- Test explicitly for crisis-state usability. Users in acute legal distress make different decisions than users in calm conditions; the interface must account for this.

Data & Privacy

- Adopt data minimization principles. Collect only what is necessary for the legal task at hand.
- Avoid retaining personally identifiable legal information beyond the active session unless affirmative, informed consent is obtained.
- Prefer self-hosted or locally deployed models for tools serving immigration, domestic violence, or other populations for whom data exposure carries physical safety risk.

Oversight & Accountability

- Require ongoing human review of a statistically meaningful sample of outputs, with documented error-correction protocols
- Establish adverse outcome tracking. Build mechanisms to identify when tool outputs correlate with case losses or harmful user decisions.
- Conduct regular bias audits disaggregated by race, language, housing type, and immigration status.
- Maintain version control and rollback capability so that model updates that degrade performance for specific populations can be identified and reversed.

III. **Guidance for verification processes and systems to ensure that AI systems provide self-represented parties with accurate information**.

Legal Content Verification at the Source

- Maintain a curated, jurisdiction-specific legal knowledge base built from verified primary sources, including statutes, regulations, court rules, and binding case law, rather than relying on general web-trained models to surface legal content.
- Establish a formal content review protocol in which licensed attorneys review and approve all legal information before it is incorporated into the system's knowledge base, with documented sign-off and version tracking.
- Implement a system that flags statutory and regulatory changes, new case law, and court rule amendments within defined jurisdictions, triggering mandatory human review before updated content is deployed.
- Never allow the system to generate legal content from parametric model knowledge alone. All legal claims should be traceable to a verified, dated, retrievable source document.
- Build retrieval-augmented generation (RAG) architecture with source pinning so that outputs are grounded in specific documents rather than probabilistic text generation, and so that citations can be independently verified.
- Require review for stale content, so unreviewed content expires rather than lingering indefinitely

Output Verification

- Require the system to display the source document underlying every legal claim it makes, with direct links to the primary source wherever technically feasible.
- Implement automated citation verification that confirms every case name, statute citation, and regulatory reference exists and says what the system claims before the output is delivered to the user.
- Build confidence scoring into outputs, with plain-language disclosure when the system is operating in areas of lower certainty such as unsettled law, sparse training data, or queries near the boundary of the system's defined scope.
- Flag legal claims that depend on judicial interpretation rather than clear statutory text, and disclose that outcomes may vary by court, judge, or jurisdiction.
- Prohibit outputs that contain unattributed legal conclusions — every substantive legal statement should be anchored to a source the user can inspect.
- Verification protocols must be realistic for lay users and should not assume access to Westlaw, legal training, or the ability to parse case law independently.
- Where feasible, the tool should show its work by displaying the rule, form, or source document on which a conclusion rests.
- Verification burdens should be shared between system design and user practice; the tool should do as much source-checking as possible before placing that burden on the litigant.

Human-in-the-Loop Review Processes

- Establish a structured sampling protocol in which licensed attorneys review a statistically meaningful percentage of system outputs on a regular cycle, with higher sampling rates for high-stakes query types (eviction, garnishment, protective orders).
- Create a flagged query queue for human review when the system detects that a user's situation involves unusual complexity, multi-jurisdictional issues, or fact patterns outside its core training scope.
- Build user-initiated escalation pathways that allow a self-represented party to request human review of any output without having to explain why.
- Implement post-interaction attorney spot-checks in which a sample of completed sessions is reviewed for accuracy, harmful omissions, and inappropriate scope creep.
- Document all human review findings in a centralized error registry that informs ongoing system improvement and tracks whether identified errors have been corrected.

Adverse Outcome Monitoring

- Develop case outcome tracking mechanisms wherever feasible, linking system interactions to downstream court outcomes to identify whether specific output patterns correlate with case losses or harmful decisions.
- Build user feedback loops that allow self-represented parties to report when information they received turned out to be wrong or unhelpful, with responses reviewed by a human rather than processed algorithmically.

- Conduct structured post-deployment audits at defined intervals assessing accuracy, bias, and scope compliance across the full range of query types the system has encountered in practice.
- Track and analyze error patterns disaggregated by query type, jurisdiction, and user population to identify whether the system performs differently for different communities. Treat differential performance as a correctable defect, not an acceptable variance.
- Establish a mandatory incident response protocol for identified errors that caused or could have caused material harm to a user, including root cause analysis, system correction, and outreach to affected users.

Structural and Institutional Accountability

- Designate a named responsible attorney or legal oversight officer with documented authority to halt deployment of content or features that fail verification standards.
- Require third-party accuracy audits by independent legal and technical reviewers before major system updates are released, particularly those that expand the system's jurisdictional scope or subject matter coverage.
- Maintain immutable output logs that preserve a record of what the system told each user, enabling reconstruction of outputs for purposes of error review and accountability.
- Publish transparency reports at regular intervals disclosing error rates, correction timelines, audit findings, and scope limitations in plain language accessible to the communities the system serves.

IV. In light of the fact that consumers have "path of least resistance" alternatives to bespoke, reliable, and trusted systems, (e.g., ChatGPT, Claude, Gemini) how do we provide incentives that will steer them to the systems that will provide the greatest benefit?

Make the Trusted Tool the Easier Tool

- Reduce friction relentlessly. A bespoke tool that requires account creation, lengthy intake forms, or confusing navigation will lose to ChatGPT every time, regardless of accuracy superiority; the interface must be as fast and intuitive as general-purpose alternatives.
- Meet users where they are, meaning trusted tools should be embedded directly in the platforms self-represented parties already use: court websites, legal aid portals, 211

systems, tenant union websites, and public library interfaces, so the bespoke tool is the first thing they encounter rather than something they have to seek out.
- Optimize for mobile-first access. The populations most likely to face eviction or debt collection are disproportionately likely to be accessing legal help from a phone, not a desktop.
- Design for zero-learning-curve entry. Users should be able to describe their situation in plain, conversational language and receive immediate, useful output without having to understand how the system works.

Build Trust Visibly and Explicitly

- Earn and display institutional endorsements from courts, legal aid organizations, bar associations, and the access to justice commission. Users are more likely to trust a tool that carries the visible imprimatur of institutions they already recognize as authoritative.
- Be transparent about what the tool does and doesn't do in plain language.
- Highlight the human backstop. If the tool connects to real attorneys or legal aid staff for escalation, make that visible; the knowledge that a human is in the loop is a meaningful trust signal for users navigating high-stakes situations.
- Use authentic community voices in outreach and interface design. Testimonials and guidance from people who look like, speak like, and share experiences with the intended user population are more persuasive than institutional messaging.

Leverage Trusted Intermediaries

- Train and equip court SHERLOCK staff, social workers, librarians, and community health workers to actively refer clients to bespoke tools rather than leaving them to find their own way to general-purpose AI.
- Partner with tenant unions, community organizations, and faith communities to build awareness and normalize use of trusted tools within networks that already have credibility with target populations.
- Integrate referrals into existing crisis touchpoints, including eviction notices, debt collection letters, and utility shutoff warnings are moments of acute legal need; partner with courts and agencies to include trusted tool referrals in the documents people receive at those moments.
- Engage public defenders, social services agencies, and benefits navigators as active referral partners, recognizing that legal problems rarely arrive in isolation.

Policy and Structural Incentives

- Get court endorsement or preferential recognition of the verified, attorney-supervised tools. If a court self-help center or judge signals awareness of a trusted tool, that carries weight with self-represented parties navigating the same system.
- Work with bar associations and legal aid funders to establish certification or quality standards for AI legal tools, so that "certified" status becomes a meaningful signal users and intermediaries can rely on.

- Seek funding that allows bespoke tools to be offered at no cost, removing any economic incentive to default to free general-purpose alternatives.
- Explore liability frameworks that create consequences — for platforms, not users — when general-purpose AI tools cause demonstrable harm through inaccurate legal guidance, creating market pressure toward accuracy and appropriate scope.
- Guidance for general-purpose tools should focus on safer uses, clear limits, and accessible verification steps rather than on formal or de facto prohibition.

Publicize Outcomes

- Invest in outcome research that documents, with rigor, the difference in case outcomes between users of trusted bespoke tools and users of general-purpose AI.
- Publish accessible case studies that illustrate in concrete, human terms what accurate versus inaccurate AI legal guidance means for a real tenant or debtor. Storytelling can be very powerful.
- Share findings openly with the access to justice community rather than treating outcome data as proprietary, building collective evidence that trusted tools produce measurably better results.

IV. How do we incorporate the systems Colorado already has in place (Sherlocks in courthouses, Law Help Colorado) to help self-represented parties use AI systems.

See question IV above.

- Integrate AI tools with Colorado's existing support infrastructure, including Sherlocks, LawHelp Colorado, self-help centers, and law school clinics, rather than creating a parallel system.

- When a user's situation exceeds the tool's reliable capacity, the tool should actively facilitate a connection to an appropriate human resource (e.g., a Sherlock, legal aid attorney, or clinic) rather than issuing a disclaimer and leaving the user to navigate on their own. A disclaimer is not an effective referral.

V. Should, and to what extent, should a self-represented party's communications with AI systems be privileged?

- The privilege question should be flagged as unsettled and appropriate for explicit treatment by the Colorado Supreme Court or legislature rather than left to ad hoc case-by-case improvisation.

VI. What safeguards prevent AI legal tools from deepening the existing two-tier justice system, rather than bridging it?

- Avoid layering on so many warnings, verification hurdles, or mandatory steps that only sophisticated users can realistically benefit from the tool.

- Describing what courts and developers should not do is not the same as closing the justice gap. We should consider crafting a guidance document that specifies affirmative obligations (e.g., what should be built, funded, and made available) rather than simply promulgating prohibitions.

Design

- Be clear about who the tool is built for; tools serving self-represented parties must be designed from the outset with the lowest-resource, highest-need user as the primary design subject.
- Require that user research and usability testing be conducted with actual members of the intended population: low-income tenants, non-English speakers, people with limited digital literacy. Testing must take place before deployment, not after.
- Optimize for mobile-first, low-bandwidth access rather than assuming broadband and desktop environments
- Mandate multilingual functionality as a core feature, not an accessibility add-on.
- Design interfaces that function for users with low literacy, cognitive load from crisis, and limited prior legal knowledge.

Content

- Ensure that the legal knowledge base reflects the rights available to low-income parties, including tenant rights, consumer protections, public benefits, debt collection.
- Audit training data and knowledge bases explicitly for coverage gaps in consumer, housing, family, and immigration law, the practice areas most relevant to self-represented parties.

- Build tools that affirmatively surface rights and defenses that self-represented parties may not know to ask about, rather than responding only to queries. Systems should not just be reactive.
- Ensure that procedural knowledge, such as how to file, where to file, what deadlines apply, what forms to use is as robust as substantive legal knowledge, because procedural failures are among the most common and consequential ways self-represented parties lose cases they should win.

Access

- Advocate for and build toward free access for self-represented parties. Tools that require subscription fees, premium tiers, or paid upgrades replicate the economic stratification of the legal market they are meant to disrupt.
- Embed trusted tools in publicly funded and community-based access points, including courts, libraries, legal aid offices, social service agencies, community health centers, rather than relying on users to find them independently through search engines that favor commercial results.
- Address the digital divide directly by partnering with libraries, community organizations, and legal aid providers to offer assisted access for users who lack devices, connectivity, or digital confidence.
- Recognize that language, disability, and digital literacy barriers are defining features of the population most in need.

Performance

- Require disaggregated performance metrics that measure accuracy, usability, task completion, and case outcomes broken down by race, language, income, housing type, and immigration status. Aggregate accuracy statistics that mask differential performance by population are a tool of the two-tier system, not a remedy for it.
- Treat differential performance as a design defect, not a population characteristic. Be sure that the tool doesn't work better for English-speaking users than Spanish-speaking users, for example.
- Conduct outcome-linked audits that assess whether the communities the tool was designed to serve are achieving better legal outcomes as a result. Process metrics alone are insufficient if they are not connected to whether people are winning cases they should win, asserting rights they are entitled to, and avoiding harms they should be protected from.
- Publish disaggregated performance data publicly and regularly, creating accountability to the communities served rather than only to funders and developers

Institutional Safeguards

- Ensure that governance structures include meaningful representation from the communities the tool serves and not merely as advisory input but as decision-making

authority over design priorities, deployment decisions, and remediation of identified harms

- Structure funding and organizational accountability to create durable commitments to underserved populations. Tools housed in nonprofit or public interest organizations with explicit equity mandates are structurally more resistant to mission drift than commercially operated tools with access to justice as a secondary use case.
- Require that legal aid organizations, public defenders, and community advocates have ongoing, institutionalized roles in tool governance, not merely as launch partners but as permanent stakeholders with authority to raise concerns and require responses.