

Guide for Evaluating AI Vendors: Key Issues and Questions for Legal Aid Programs

Authored by: Ashurst Perkins Coie LLP and the Lawyers Trust Fund of Illinois

This resource provides guidance for legal aid organizations seeking to utilize AI technology and tools. It provides a high-level set of issues and questions to consider when evaluating vendors that provide AI tools to help ensure the chosen solutions align with the organization's legal needs, standards, operational requirements, and mission.

This guide is a supplement to the organization's standard diligence process for onboarding new vendors. For this reason, standard diligence issues are not addressed, such as pricing, customization and integration support, Service Level Agreements ("SLAs", which define performance expectations like uptime guarantees), or other key contract terms.

Note: It is the nature of AI tools that users can enter or upload pretty much anything as a prompt. As a result, there may be issues of potential misuse of client data or the organization's own confidential information even if the tool is not designed to act on that type of data. For example, a user could upload a client's health records or veterans' benefits termination notice or other supporting documents to a publicly available general purpose AI tool (*e.g.*, ChatGPT or Gemini) and ask the AI tool for help summarizing the document or drafting an appeal letter. The result is that the document could become part of the AI's training data. In addition to asking the below questions, it is vital that organizations have AI use policies and users are trained on which tools to use for what purposes and how to protect sensitive information when using more open tools.

1. Functionality

<i>Questions</i>	<i>Comments</i>
What are the intended uses and limitations of the AI tool (e.g., general purpose, legal research, drafting, translating)?	Based on the AI tool's intended purpose(s), you should consider whether this product is fit for your intended use. For example, some AI tools, like ChatGPT, are general purpose tools, while others are developed with a very specific use in mind such as legal research or human resources management. Ways to determine if a tool is fit for your intended use: • Review the tool's documentation and specifications for its stated purpose and performance claims; • Consult with other users or third-party reviewers who have experience with the tool; or • Request a demo from the vendor.

<i>Questions</i>	<i>Comments</i>
Does the AI tool offer different versions or add-on features? If so, what are they? Do they cost extra?	Different versions of a tool or add-on features may offer varying levels of functionality, accuracy, performance, support, or more user controls. Examples of additional versions to consider: • Free (typically has the fewest features, usage limitations, and little support) • Standard/Professional (typically includes core features, moderate usage limitations, standard support) • Industry specific (e.g., healthcare, finance, legal specific features) • Non-profit specific (discounted access) • Enterprise (advanced customization for a specific organization) Examples of add-ons to look for, include: • Language support (multilingual interfaces, translation support) • Enhanced customer support (24/7 support, faster response time; onboarding assistance) • Privacy controls (opting out of model using your personal data for training purposes)
How does the AI tool integrate with our internal systems and workflows? What mechanisms are available (e.g. APIs, connectors, SDKs) for ingesting or acting on internal data securely and reliably? Is integration support provided?	Understanding how the tool integrates with your internal systems is important for ensuring that it can seamlessly connect with your existing workflows. APIs (tools that let different software systems communicate) enable seamless connection between the AI tool and your existing platforms (e.g., CMS, Microsoft 365, Google, etc.), allowing for automation and customization. Strong vendor support ensures smooth implementation and reduces technical challenges. Software development kits (SDKs) are prebuilt interfaces and tools that you can use to make applications and systems quickly with as little code as possible. Support may include technical guides on APIs, SDKs, and configuration steps, implementation and

<i>Questions</i>	<i>Comments</i>
	ongoing technical support, and API monitoring tools to track performance.
What type of onboarding, training, and customer support is provided?	Many AI tools require training to ensure they are used appropriately. Compare the level of onboarding support and training provided by vendors to the amount you anticipate needing for your employees and volunteers.
What assurances are provided about the future financial sustainability of the AI tool?	When assessing the financial sustainability of the AI tool, look for indicators such as years in business, customer base, annual revenue, amount of external funding raised, and ongoing support commitments by the vendor.
Is the user interface of the AI tool accessible and compliant with Section 508 of the Rehabilitation Act of 1973 and/or Web Content Accessibility Guidelines (WCAG) 2.1? If so, what is the level of compliance?	Section 508 requires that federal agencies' electronic and information technology is accessible to people with disabilities. WCAG 2.1 are internationally recognized standards for web accessibility, outlining specific criteria to ensure digital content is accessible to individuals with a wide range of disabilities. When evaluating vendor responses, look for: • Clear statements of compliance • Supporting documentation, such as a Voluntary Product Accessibility Template (VPAT) • Details about the specific level of WCAG conformance (e.g., Level A, AA, or AAA) • Vendor's ongoing accessibility testing and support for users with assistive technologies • Support for accessibility issues

2. AI Tool Development & Training

<i>Questions</i>	<i>Comments</i>
What data was used to train the AI tool? How are the data sources, methods, and algorithms that the AI tool uses to perform tasks validated, tested, and updated? How frequently is the training data reviewed or audited?	If a vendor is unwilling to provide information about the origin, collection methods, and processing history of the data used to train and operate the AI tool, that may signal that the tool was trained on unlicensed or scraped data. Understanding the sources of training data will also help identify potential biases, inaccuracies, or limitations with the AI tool. For example, AI tools trained on large amounts of data and content scraped from the Internet may be more prone to hallucination compared to an AI tool trained on a more select, curated set of materials such as collections of case law or legal treatises. Understanding sources will also help determine what representations, warranties and indemnities to seek in your contract with the vendor. Examples to look for: • Representations that the data is lawfully sourced and used in compliance with applicable laws or that reasonable efforts have been made to identify and mitigate bias in the training data • Warranties that the tool is trained on accurate data • Indemnities protecting you from third-party suits (IP infringement or unauthorized use of data) Closely evaluate how often the vendor reviews or audits its training data. If the vendor does not conduct these assessments at least annually, there may be an increased risk of the AI tool producing inaccurate or irrelevant/outdated results. (Note, the appropriate frequency of review may be determined based on the level of risk associated with the AI tool).

<i>Questions</i>	<i>Comments</i>
Who was primarily responsible for developing and training the AI tool? Was there collaboration with a third party? If yes, who is the third party and at what stage(s) was it involved (e.g., planning, design, development, testing, deployment, etc.)?	Confirming whether the AI tool uses a model developed by a third party can help you identify and assess AI tool-related risks, such as ongoing litigation related to the AI tool, potential service disruptions, or dependencies on third-party updates. Many specific purpose AI tools are built on top of large language models (LLMs) developed by third parties. For example, Harvey AI, an AI legal assistant for law firms, is built on OpenAI's LLM. Understanding the AI tools underlying model and its developer can be helpful in evaluating legal and operational risks related to the tool.
Are user inputs (information, prompts, or materials provided to the AI tool by the user) used to train or refine the AI tool? Can users opt-out of having their data used for training purposes?	Be very cautious about allowing vendors to train or refine their AI tools based on data that you provide through use of such tools. Depending on the use and the types of data the AI tool processes, this practice could result in your confidential or proprietary information being compromised and even disclosed to other users of the AI tool. Many vendors offer free versions of their AI tool that may, by default, use the data you put into the tool (e.g., what you type in, prompts, uploads) and the resulting information for training purposes. In contrast, paid or enterprise versions often provide you with enhanced privacy controls and greater flexibility over how your data is used. If a vendor offers you the ability to opt-out of having your data used for training or refining purposes, it is strongly recommended that you exercise this option and ensure that individual users cannot override such opt-out via their account settings.

<i>Questions</i>	<i>Comments</i>
Where is the AI tool deployed? Who hosts the AI tool?	Deployment options may include cloud (e.g., on the internet) or on-premise (e.g., on your company's own computers) or a hybrid approach. Understanding where the software is deployed can help you understand the technical and operational requirements and implications of using the AI tool, and the potential risks and benefits of the different deployment options. For example, cloud solutions can be easier to set up and can integrate well with other tools or programs such as Microsoft 365, but you may have less control over where your data is stored. On-premise tools give you more control over your data but can be more expensive and harder to maintain and expand as your needs grow.
How is the AI tool maintained?	Confirm what ongoing support and services the vendor provides to maintain the AI tool. Examples include: • Regular software updates • Compliance updates • Help desk support • System monitoring and maintenance (e.g., identifying and fixing bugs/issues)

3. Inaccuracy, Discrimination & Bias	
Questions	Comments
How is the quality, accuracy, and reliability of the training data and the results it produces ensured?	Evaluating the mechanisms a vendor has in place to continuously monitor and improve their AI tool is crucial for ensuring that the training data and results produced by the tool remain accurate, reliable, and as unbiased as possible. Examples of such mechanisms include: • Regular testing of results to identify and address errors/inconsistencies • Performance audits to assess the tool's effectiveness • Bias detection tools • Human oversight (e.g., expert review of results produced by AI tool) • User feedback loops (e.g., enabling users to report issues and investigating the issues) Vendors that cannot speak to such measures may be less sophisticated or may be less capable of ensuring the AI tool's quality, accuracy and reliability.
How are ethical considerations and potential biases in AI tools addressed? What steps are taken to mitigate the risk of the AI Tool generating results that are biased or discriminatory?	AI tools often inadvertently generate outcomes and results that can be biased or discriminatory. Vendors should have processes in place for identifying and mitigating bias to minimize the risk of biased or discriminatory results. Such processes may include: • Conducting routine bias audits or impact assessments (and making results accessible to stakeholders) • Using training datasets that reflect the diversity of populations served • Providing ongoing training for developers on recognizing and mitigating bias
Does the AI tool document or provide sources for its outputs/responses (results,	Some AI tools do not “footnote” the sources relied on for prompt replies, meaning there is no easy way to confirm the accuracy of results.

content, or materials produced by the AI tool in response to a user's prompt)?	Understanding whether the vendor provides sources for its responses can enable you to verify the credibility and traceability of such responses and understand how the AI tool supports your due diligence and verification processes.
How is the AI tool programmed to respond if it does not know or have an answer?	Be cautious about AI tools that do not clearly indicate when a question or task falls outside of its knowledge base or capabilities. If the tool attempts to generate a response in such situations, it may provide inaccurate or misleading information, which can create risks for both users and clients (e.g., reputational damage to the organization, incorrect legal advice to client).
Has the AI tool been trained with data from diverse populations, especially marginalized groups?	Evaluate the fairness and inclusiveness of the data sets used to train the AI solution to ensure that the tool's responses are accurate and relevant for your clients.
Do customers have a way to flag when the output/response is incorrect, shows bias, or has other issues? How is that feedback communicated back to the vendor?	You should not assume that responses generated by an AI tool are accurate or free from bias. Vendors that provide mechanisms to report issues and welcome feedback from users, experts, and the broader community to help refine and improve their practices and technologies may be more sensitive to and on guard against these risks. Such vendors are also better positioned to address these risks because they proactively collect instances of inaccurate or biased responses in real-time across their customer base.
Does the AI tool help make decisions?	Decision-making AI tools that include limited or no human involvement can pose a greater risk of potential biases, ethics concerns, and inaccuracies.

4. Compliance, Privacy and Data Security

<i>Questions</i>	<i>Comments</i>
Does the vendor have an AI governance program? If yes, how does the program evolve with emerging AI regulations, security frameworks, and/or a code of conduct?	An AI governance program is a structured framework of policies, processes, and controls that ensures the ethical, responsible, and compliant development and use of AI within an organization. If the vendor has a program, discuss the program to help ensure the vendor's ongoing compliance with relevant legal and ethical standards and best practices for developing and deploying AI solutions.
What privacy laws does the AI tool comply with and what privacy certifications does the vendor have?	Compliance with recognized privacy standards and privacy certifications can serve as indicators of strong security and privacy practices. Examples include: • Children's Online Privacy Protection Act (COPPA) • Federal Trade Commission (FTC) Act • NIST Privacy Framework • SOC 2 Type II • ISO/IEC 27001 • FedRAMP
What type of information do you intend to input into the AI tool (e.g., text, image, audio, video)? Will that information be confidential, attorney-client privileged, health data, or financial data?* <i>*This is an internal question for the legal aid organization.</i> If the information is confidential, privileged, or health or financial data then ask the vendor – what safeguards and remedies are in place to help ensure the privacy, security, and confidentiality of this data?	Examples of safeguards and remedies to look for: • Data encryption • Strict access controls over who can access the data • Opt-out and data deletion mechanisms • Zero data retention agreements Note, if inputting information that is confidential, attorney-client privileged, health data, or financial data, you should refrain from using free or less-protective versions of an AI tool. Under these versions, the AI vendor may be able to store, review, or use the information you input to train future AI models.

<i>Questions</i>	<i>Comments</i>
Does the AI tool process “personal data” (or “personal information”), in either the training phase or when in live use? Note that such terms are defined very broadly under U.S. privacy laws and include unique identifiers that can be used to recognize a consumer, household or device, such as device identifiers, IP addresses, cookies, beacons, pixel tags, mobile ad identifiers, customer IDs, and user aliases. If yes, can materials documenting the vendor’s compliance with applicable data protection laws, regulations and best practices be provided?	Press the vendor on how it defined “personal information” in light of its answer. For example, if the vendor takes the position that the AI tool only processes “de-identified” or “aggregated” data, which are both defined terms under U.S. state privacy laws with specific requirements, you should ask the vendor to demonstrate how they satisfy these requirements.
If the AI tool processes personal data, does the AI vendor act as a “processor” or “service provider”? If considered a processor, does the vendor have a data processing agreement? Is the vendor a “controller” for regulatory purposes?	Vendors that contractually agree, through data processing agreements, to act as “processors” or “service providers” are restricted from processing your data for purposes other than to provide the requested product/service to you and for limited internal purposes. Accordingly, such designation can help you ensure that the vendor does not process the data it receives from you for the benefit of other customers or for other commercial purposes. If a vendor is the “controller” similar restrictions may not apply by default. As a controller the vendor may set its own purposes and legal bases for processing and may use your data for its independent activities (e.g., monitoring and analyzing trends, usage or product improvement) subject to applicable law and its privacy notices, unless you contractually limit these uses (which may be limited based on the specific version of the AI tool being used and the size of the vendor). However, the likelihood that an AI vendor will be considered a controller is low. Regardless of the vendor’s designation, however, you should carefully evaluate the rights vendor reserves with respect to your data. For instance,

<i>Questions</i>	<i>Comments</i>
	does the vendor reserve rights to train its algorithms using personal data collected from you? If yes, then the vendor’s activities may fall outside the scope of what is permissible for a “service provider” or “processor” and you would need to decide if you are comfortable with this approach and address related privacy compliance obligations (e.g., related to sales of data under state privacy laws).
What sorts of privacy controls are available to customers?	Most AI vendors offer privacy controls to protect personal information and ensure compliance with data protection laws. Access to these controls may vary based on the version of the tool you have or the controls may be add-on features. Some examples of privacy controls include: • Opt-out mechanisms – gives user the option to prevent their data from being used to train or improve AI models. • Data retention and deletion controls – allows users to set how long their data is retained and to request deletion of data after processing • “Private” or “stealth mode” – keeps a user’s inputs/prompts and the AI tool’s outputs/responses private (e.g. chat history will be disabled) • Data anonymization and masking – tools that remove or mask personal identifiers from datasets before processing (e.g., name, SSN) The extent to which a vendor provides such features may indicate the level of sophistication of the vendor, as these topics are common concerns of enterprise customers. For vendors that do not offer any such controls, you should carefully consider whether the benefits of the vendor’s AI tool outweigh the lack of ability to take advantage of privacy, stealth mode, and retention controls, as such controls are advantageous to you for customization, oversight, and privacy preservation purposes.

<i>Questions</i>	<i>Comments</i>
Where is personal data processed and stored?	Confirm the region(s) in which the vendor processes and stores any personal data that you may provide through your use of the AI tool, as some jurisdictions pose a higher risk in the event of government requests to access data. Depending on the sensitivity of the data processed by the AI tool, you should consider seeking assurances from the vendor that data will not be processed or stored in high-risk regions, e.g., areas where there may be a lack of comprehensive data protection laws or increased government surveillance or interception powers.
Does the vendor have data retention policies in place? If so, what are they? More specifically, can users control how long data is retained or if it is retained? What happens to a user's data after the retention period is over?	Data retention policies are essential for ensuring that personal and sensitive information are managed responsibly. Note, the ability to control how long data is retained or if it is retained may depend on the version of the AI tool (e.g., only ChatGPT Enterprise, ChatGPT Edu, and ChatGPT Team enable users to control how long their data is retained). Additionally, if required by law (e.g., for investigative or litigation purposes), vendors may be required to retain your data even if they have policies in place that limit data retention to a certain time period.
Has the vendor prepared a data protection impact assessment/AI risk assessment? Can the vendor share the assessment with you for your review? To the extent your organization is subject to assessment or audit requirements related to its use of the AI tool, will the vendor assist/cooperate with you to satisfy these requirements?	In many cases, privacy laws require companies to conduct a data protection impact assessment prior to developing or deploying AI tools. Although there is no active federal AI law, several states, including California, Colorado, New York, Utah, and Texas, have enacted or advanced AI legislation that, in certain high-risk contexts (e.g., employment, financial or lending services, or housing) or consumer-facing contexts, require companies that develop or deploy AI tools to prepare impact or risk assessments of such tools and implement safeguards to protect against discriminatory or deceptive outcomes.

<i>Questions</i>	<i>Comments</i>
What sort of indemnities does the vendor commit to providing?	Confirm whether the vendor offers their customers indemnities for issues, such as IP infringement claims related to the outputs/responses of the AI tool, data breaches or other security incidents, or general operational failures (e.g., errors in AI responses or decision-making). Such offer suggests a higher degree of confidence in the AI tool and the responses it produces.

5. Intellectual Property Considerations

<i>Questions</i>	<i>Comments</i>
Who owns the inputs/prompts and the outputs/responses? What rights does the vendor reserve to use such inputs/prompts and/or outputs/responses generated by the AI tool?	It is rare for vendors to insist on owning the outputs/responses of their AI tool, and even more rare for vendors to insist on owning a customer's inputs/prompts to their AI tool. However, if there is a compelling reason for doing so, or if a vendor reserves any rights to inputs/prompts or outputs/responses, you should ensure that the AI tool is limited to specific, pre-approved users and that employees are prohibited from using the AI tool to provide such prompts or generate responses that incorporate personal, confidential, or proprietary information, or key or important creative assets.
How do you plan to use the outputs/responses produced by vendor's AI tool?* <i>*This is an internal question for the legal aid organization.</i>	AI tools whose responses will be used externally (e.g., tools to generate marketing copy) may involve higher practical risks than AI tools whose responses will only be used internally. For example, if a tool's responses will be used externally, be attuned to risks of copyright infringement or improper use of images or an individual's name, photograph, or likeness. Copyright law protects original works such as text and images, and using these without proper authorization can result in liability, even if the content was generated by an AI tool. Similarly, using a person's name, image, and likeness without their consent can violate privacy and publicity rights, which are protected under state, and sometimes, federal law.

This guide was developed by Ashurst Perkins Coie and the Lawyers Trust Fund of Illinois as general information and not as legal advice. No attorney-client relationship is created with the reader. The information provided is not intended to be a comprehensive review of all developments in the law and practice, or to cover all aspects of those referred to. Readers should take legal advice before applying it to specific issues or transactions.