

Towards Human-Centered Standards for Legal Help AI

Margaret Hagan, Stanford Law School

October 2023

Abstract

As more groups consider how AI may be used in the legal sector, this paper envisions how companies and policymakers can prioritize the perspective of community members as they design AI and policies around it. It presents findings of structured interviews and design sessions with community members, in which they were asked about whether, how, and why they would use AI tools powered by large language models to respond to legal problems like receiving an eviction notice. The respondents reviewed options for simple versus complex interfaces for AI tools, and expressed how they would want to engage with an AI tool to resolve a legal problem. These empirical findings provide directions that can counterbalance legal domain experts' proposals about the public interest around AI, as expressed by attorneys, court officials, advocates, and regulators. By hearing directly from community members about how they want to use AI for civil justice tasks, what risks concern them, and the value they would find in different kinds of AI tools, this research can ensure that people's points of view are understood and prioritized, rather than only domain experts' assertions about people's needs and preferences around legal help AI.

Keywords

Legal technology, Artificial Intelligence, Participatory policy-making, Legal Design, Access to Justice

Funding

This research was done with no external funding, only department resources.

Acknowledgments

Nora Al Haider

Introduction

With the introduction of new generative AI models and agents, justice system professionals are cautiously excited about how AI might improve people's ability to get and use legal help. In conversations among lawyers, court staff, and academics who work on the civil justice system, these domain experts discuss how AI's advances might close the justice gap for people facing common problems like evictions, poor living conditions, debt collection lawsuits, divorce, and custody matters.^{1 2} Legal help AI might help a person spot that they have a legal issue, triage them to appropriate resources, guide them to free assistance and help them present and negotiate their case in order to get to a good resolution.³

Along with these optimistic discussions of AI's potential, domain experts also warn about how AI tools could pose risks and harms to the public, as well as to the justice system.^{4 5 6 7} Might the AI give people incorrect legal information that they rely on when making high-stakes decisions?⁸ Might AI tools be a second-class service, that those without resources use rather than supposedly higher-quality human services?^{4 5} Or might people not be able to use AI tools because they don't have reliable Internet access or technical knowledge to make use of them?⁴

This debate about techno-optimism versus -skepticism about legal help AI is largely speculative. Those debating opportunities, risks, harms, and protections are domain experts (rather than community members), proposing possible versions of how people might access and use AI for legal problems, or how they might not access and how they might not use it.^{1 2 3 4 5 6 7 8} The optimistic and skeptical domain experts both frame their discussion around the public interest and base their positions on their assumptions of the public's behavior, capability, and preferences. Because generative AI is so new, there are few applications in the legal help domain, and domain experts have little beyond their own speculation to determine the actual opportunities and risks for legal help AI.

This article balances this expert-driven approach with a community-driven one. When deciding what kinds of AI tools should and should not be developed, and what kinds of technical limitations, interface designs, and policies should be defined, experts' speculation about the public interest should not be the only input. Rather, experts' preferences and predictions about the public's use of AI for legal help should be counterbalanced with data about the actual preferences and behavior of the public, especially community members who need legal help and would be in the potential audience for legal services. As policy governing AI in the legal domain is developed, community members' perspectives should be included to ensure that these policies match the needs and realities of the community.^{9 10}

This article proceeds as follows. It reviews the burgeoning research and policy discussions about how AI can help or undermine community members' access to the civil justice system. Then it discusses the importance of participatory design research, which can ensure that technology development, investment, and policy-making is shaped by the potential users of the new technology, and not just professional experts.¹¹

The article then presents the survey and design session used to conduct exploratory, qualitative discussions with community members about whether they would use generative AI tools for legal help problem-solving, how they would behave, and what kind of interface, complexity, and guardrail designs they would prefer. The article presents the results from the interview sessions, highlights key themes and

insights that emerged, and concludes with the next steps for research and policy-making, as well as hypotheses to explore.

This article offers empirical substance to the policy discussions about what people want from AI when trying to understand and resolve legal problems. Its initial qualitative research findings can be a blueprint for further extensive and ongoing research efforts that can inform the development of AI for legal help use cases, as well as the policies that determine what should be built, what behaviors and interactions should be guarded against, and what accountability mechanisms should exist to mitigate harms or punish violations. This initial research cannot fully inform development and regulation, but it lays out initial findings and research protocols that can expand to include more voices and gather input on forthcoming legal help AI models and applications.

Background on Legal Help AI Policymaking

Throughout their lives, individuals across the population encounter legal problems, or ‘justiciable’ problems, which affect their financial, housing, and family stability, as well as their personal security.^{12 13} These problems include situations like debt collection, poor living conditions in their rental housing, evictions or foreclosures from their home, harassment or wage theft at their job, child custody and support issues, or difficulty accessing government benefits.¹⁴ Many people do not, or are not able, to resolve these problems using the justice system.^{15 16 12 13 14} They may not realize that this life problem has a legal dimension, and thus that a lawyer, court staff member, legal help website, or another help service could help them understand their options, make a plan, and come to a resolution.^{17 15} Or, they may try to reach out for help by calling a legal aid group, filling in an online intake form at a legal website, or showing up to a court help center, but then either be turned away because of a lack of capacity or only receive a 30-minute service rather than substantial and ongoing help.¹³

As new technologies have emerged over the past several decades, domain experts have spotlighted how they might improve people’s ability to do legal tasks. Some of these tasks include spotting legal issues, understanding their options, finding free services, filling in important paperwork, preparing for hearings and negotiations, and coming to final resolutions.^{18 19 20 21 3} With growing awareness of generative AI, more domain experts have identified ways that AI might help people access justice and get legal help.^{22 23 24 25} In this techno-optimistic way of looking at generative AI, domain experts imagine that people will find and use these technologies to improve their capacity to address their legal problems, access their rights and resources, and improve their life outcomes.²⁶

Techno-skepticism has emerged in response, with other legal domain experts imagining future consumer harms, poor justice outcomes, or worsened inequality if more AI was developed to help people deal with their legal problems. One common concern is that AI might produce hallucinations about legal rights, procedures, and rules that people rely on to their detriment -- e.g., losing a case, missing a key deadline, or having a claim rejected.²⁷ Another speculative concern is that legal help AI might be developed in inequitable ways, with only people who have strong English and technology literacy able to benefit from it.^{4 5 6 7}

This growing debate between domain experts foreshadows more explicit policymaking around legal help AI in the near future. The discussions that now are happening in academic journals, industry conferences webinars, and trade publications will likely soon harden into policies that govern what kinds of legal help

AI is developed, what guardrails and limitations are built into models and applications, and what consequences are faced by those who offer or use it. These policies will likely be made by bar associations and courts that formally regulate the legal profession,^{28 29} professional associations that issue principles and frameworks that informally set out norms for legal providers,³⁰ and the technology companies that choose whether and how to offer AI models and tools for legal help use cases.^{31 32}

As this new wave of generative AI advances and expands, and new sets of formal and informal policies develop in response, there is one group that is not often present in the discussions about how AI might expand access to justice or increase legal help: the public. Apart from broad speculation about how people might want to use AI, how they would use it, and what they might get wrong with it, the discourse now is grounded in speculation about abstract community members rather than community members' voices, actions, and proposals.

In other policy domains, public participation in the development of new regulations, principles, and other policy initiatives has been recognized as a crucial way to develop more responsive, equitable, and beneficial policies.^{33 34 35} Groups working on AI development have developed initiatives to include more members of the public in the design, implementation, and regulation of AI.^{36 37 38 39} These groups have highlighted the need for more participatory policy-making as AI technology and regulation develop, which includes more consultations, testing, and collaboration with community members. Participatory policy-making shifts the focus away from domain experts and technology developers deciding what AI should be built, and what policies should govern how people use it. Community members can help set the agenda of what AI should be built, how it could be made to be most useful, and what guardrails should be established to protect people from possible risks.

In the domain of legal help AI, there has not been a shift toward participatory policy-making yet. In part because of how recently generative AI has become widely known among professionals and the public, there have not yet been consultations, academic research, design workshops, town halls, or other initiatives to hear from community members about whether or how they would want to use AI to address legal problems. Though there has been growing interest in community-centered methods to develop new legal services and improve access to justice broadly,^{40 41 42} there has not yet been a similar effort around community-centered legal help AI specifically.

The Research Study Design

In the Spring of 2023, my group at Stanford Legal Design Lab created a research protocol to initiate discussions, brainstorming, and feedback with community members about legal help AI. The group is made up of two lawyers who have worked on technology and service design projects to increase people's uptake of legal help, and we were motivated to understand if the new AI platforms may better improve outreach and connection to help resources. At the Lab, we consciously try to stay in between techno-optimism and -skepticism. Rather, we try to evaluate new technologies like AI platforms by understanding community members' perspectives about whether there is value in this technology for legal help-seeking. To this end, we created a research protocol with qualitative research methods taken from design research, participatory policymaking, and human-computer interaction.^{43 44 45} The protocol also adapts scenario-based research protocols previously used to study how people respond to life problems that have a legal dimension, and how they may attempt to use technology to understand and respond to a legal problem.⁴⁴⁴⁵

For this study, we recruited adults in America who identified as having had a civil legal problem and were open to participating in an online interview about how they use the Internet to solve problems. We used a convenience sampling method to speak to various people during this initial, exploratory research about people's approaches to generative AI, before we had any clear hypotheses about the topic. The recruitment was run through short advertisements on Craigslist and on Facebook, that directed people to a screening survey about their demographics, past civil justice problems, and their confidence in using the Internet to respond to life problems like landlord issues, problems with their employer, or a debt collection action. In June 2023, we advertised this specific study to the group of people that had signed up. The group received an email asking if they would be open to a 40-minute Zoom session about AI and legal help. Interested people were directed to a scheduling site to sign up for an interview session during June.

The online interviews were conducted by one or two research team members, using a Qualtrics survey they showed the participant through a shared screen. The lead team member asked the survey questions and typed in the participant's responses. The interview had 4 parts: (1) a review of the research study before the participant consented to continue, (2) background questions on the person's legal capability, technological capability, and experiences with AI, (3) a scenario exercise in which the person used an AI tool (in this case, Bard) to respond to an eviction notice, and (4) questions about the person's experience with the tool, feedback, and brainstorming about how to improve it, and preferences around the ideal AI tool for legal help. During the third part, the participant was given a fictional situation in which they lived in an apartment for 4 years and just arrived home to find a notice on their front door from their landlord that said they had to leave in 2 weeks or that they would be sued for eviction. The participant was then directed to share their screen, open Google Bard, and talk aloud while they used it as if they were experiencing this fictional situation.

The study was designed in order to gather people's expressed preferences about whether and how they would use an AI tool like Bard or ChatGPT to respond to a legal problem, and then also to observe their actual behavior when using it. In addition, the study was designed to engage the participant in co-design, so that they would propose changes to the interface, content, policies, or other parts of the tool or the underlying model so that the AI would better suit their needs and prevent possible harms, from their perspective. By having participants engage with the AI tool to try to solve a problem (even if it was fictional), the participants were able to give detailed, specific feedback about how the AI could be empowering and valuable to them, or how it could be harmful or irrelevant. The session could elicit more concrete feedback about the level of complexity versus simplicity that the person wanted from the AI tool. With complexity versus simplicity in regards to explanations of how the tool works (the reasons why it is giving the responses that it gives) and explanations of legal rules, procedures, and services that it mentions in its responses.

The interviews contained a mix of structured multiple-choice and slider questions, along with open-ended conversational questions. Structured questions were meant to make participants express their preferences and feelings about AI tools in defined ways. For example, participants were asked questions like "On a scale of 0 to 6, How much would you trust what an AI tool like ChatGPT would tell you, especially around the [legal] problems we've talking about?" and "Would you rather the AI tool show you explanations about why it's giving you those answers? Or that it would just tell you its answers? Please tell me where on this scale of 0 (simple) to 6 (complex) you prefer." These structured questions covered

pre-existing feelings towards generative AI and towards the legal system, then also the feedback on the specific AI tool interactions they had while trying to use it to solve the fictional legal problem, and feedback on a series of brainstormed ideas about what might improve the AI. The quantitative results that emerge from these structured questions can help categorize the different participant types and make the feedback more structured.

The open-ended conversational questions allowed the participants to express their detailed feelings and ideas about AI. For each structured question, they were given a follow-up open-ended question to gather more information about their rationale, previous experiences, ideas, and questions. The study deliberately included many open-ended questions in order to elicit as much information about people’s thought processes, preferences, trust, and feedback, because people’s relationships with generative AI is so new. We did not want to rely only on multiple choice and slider or scale questions, because they would not give us an understanding of why participants answered in the way they did. This research project is exploratory, to identify how people think about AI, what they trust, and how simple or complex they want an AI tool to be. Open-ended qualitative survey questions provide us with a clearer understanding of people’s feelings and thought processes, so we can identify common categories of AI tool users and their requirements, needs, and preferences for AI for legal help.

Participant Population

We had 15 people sign up for and participate in online sessions. All participants completed the session, and no one skipped any question or section. Of the 15 participants, 6 were male and 9 female. Seven participants identified as white, 4 as Asian, 2 as Black or African American, 2 as Spanish, Hispanic, or Latino. Participants were from California, New York, Maryland, and New Jersey. Six reported that they lived in a large-sized city, 8 in a medium-sized city, and 1 in a small city. Two were in the 18-25 age range, 2 in the 26-34 range, 9 in the 35-54, and 2 in the 55-64. One participant reported speaking English as their second language, and the remaining 14 spoke it as their primary language. One participant identified as having a disability, and the remaining 14 did not. Eleven participants reported a combined household income of \$100,000 or above; 2 reported an income in the \$90,000 range; 1 in the 60,000 range; and 1 in the \$50,000 range. Professions ranged from HR consultant, to accountant, to high school teacher, to administrative assistant, to office manager, to payroll clerk, to construction management.

See this table of the participants, with a key that will be used throughout the next section.

Participant	Age range	Education	Income	Profession	Race/Ethnicity	Sex	Disability?	Primary Language
1	26-34	Professional Degree	100,000+	Unemployed	Asian	Male	No	English
2	55-64	Masters Degree	100,000+	Management Consultant	White	Female	No	English
3	35-54	4-year college degree	100,000+	Options Trader	White	Male	No	English
4	18-25	4-year college degree	100,000+	Masters Student	Asian	Male	No	English

5	35-54	4-year college degree	90,000-99,999	Payroll clerk	Spanish-Hispanic-Latino	Female	No	English
6	35-54	4-year college degree	100,000+	Executive assistant	White	Female	No	English
7	35-54	Some College	60,000-69,999	Administrative Assistant	Asian	Female	Yes	English
8	35-54	4-year college degree	100,000+	Sales manager	White	Male	No	English
9	26-34	Some College	50,000-59,999	HR generalist	Black-African-American	Female	No	English
10	35-54	4-year college degree	100,000+	Business analyst	Asian	Female	No	Not English
11	18-25	Some College	90,000-99,999	Construction management	White	Male	No	English
12	55-64	4-year college degree	100,000+	HR consultant	White	Female	No	English
13	35-54	Masters Degree	100,000+	High school teacher	Spanish-Hispanic-Latino	Male	No	English
14	35-54	4-year college degree	100,000+	Accountant	White	Female	No	English
15	35-54	4-year college degree	100,000+	Office Manager	Black-African-American	Female	No	English

We asked participants about their confidence in being able to solve a legal problem, and their confidence in using the Internet to solve life problems. These self-assessments helped us understand our participants' legal capability and technical capability. Participants generally assessed their legal capability at 2.9 on a scale of 0-6, with a standard deviation of 1.3. Participants rated their ability to use the Internet to solve life problems at an average of 4.5, with a standard deviation of 1.4. All participants reported that they used the Internet frequently to get answers to problems they are dealing with. On a scale of 0-6 to assess the frequency of using the Internet, the average rating was 5.5, with a standard deviation of 0.7. Generally, the participants tended to be frequent Internet users, with medium to high confidence in their ability to find answers to life problems online, and slightly less confidence in their ability to address legal problems that might arise.

As this was a convenience sample that was recruited via social media ads in English, the participants do not represent the US public. There is an underrepresentation of people above 64, of lower income brackets, of limited English proficiency, of limited technological capability, and coming from rural backgrounds. Our team acknowledges this underrepresentation and recommends that future research studies (whether they be short surveys or qualitative sessions) employ sampling techniques and research protocols that ensure sufficient representation of the adult public. In our findings and discussions, we

present the results only as indications of what these specific participants did and expressed. These particular people's preferences, behaviors, and ideas can help us shape future hypotheses, research, and design efforts, but they should not be taken as representative of the broad US public.

Findings of the Interviews

This section presents the responses and behaviors of the 15 participants in the study. The subsequent section will discuss the responses, particularly the more qualitative answers, in greater depth to explore the insights, hypotheses, and patterns that the research team took away from the findings. Both sections should be read with awareness of the limited sample of people, who (while diverse in gender, race, and professional background) tended to be very familiar with using the Internet for problem-solving and tended to have household incomes around \$100,000 or more.

The findings can be grouped into three main topic areas: (1) how participants understand and think about generative AI (especially regarding life or legal problems); (2) how they interact with and assess an AI tool to address a particular legal problem; and (3) what the AI tool presents back to the participant. The findings include both the quantitative responses that participants gave when asked to rank the AI tool on scales of 0-6, as well as qualitative responses they used to explain their quantitative ranking. In this early study, our research team had only a basic coding framework, of whether the participant was positive, negative, or neutral to the AI platform. Both team members were present for the interviews, and came to a consensus on the basic coding, and also flagged certain qualitative responses that seemed to present interesting insights or hypotheses. We expect future studies to have more detailed coding, based on some of the patterns and hypotheses below.

1: How Participants Understand Generative AI

All participants demonstrated an understanding of what AI is and how it works, even if they did not agree on its trustworthiness or value. The survey asked participants to explain what ChatGPT or a similar AI tool is like, as if they were describing it to a friend who wasn't familiar with it. Participants showed an understanding about how generative AI works, with descriptions like the following.

I ask the question and it pulls information from all over the Internet to give me a specific answer (P13)

It's like a 24-7 personal assistant that you can utilize for personal and business use (P9)

People describe it as a really strange thing where we get answers for everything. People are kind of worried about it, that it will take over everything. (P6)

Like the robot in Space Odyssey HAL: it's a computer program that learns from what you and other people type, and it uses that dataset and is able to have a conversation with you. (P3)

The survey followed up with a question asking the participant to compare generative AI to something else, as in a metaphor or analogy. The most common comparison was to Google Search. Participants thought of AI tools in comparison to a search engine, often explaining that AI performs a similar task to search engines but does so in better ways.

It's an advanced form of using a Google Search. You can ask more specific questions and get more specific answers. (P12)

It's very quick and easy to ask questions, get suggestions and a detailed answer. I would compare it to Google Search. But you don't need to browse through search results. You type the question you have and it compiles & answers in 1 response. (P10)

It's like using Google but having a conversation with you.(P4)

Participants did differ in their assessment of whether AI was a good or bad thing to use. Two participants expressed reluctance to engage with AI, both based on a concern about data privacy, tracking, and future ways that the AI or the companies that run it might harm them. One of the skeptical participants (P2) explained her hesitancy around AI, saying “All the things I'm reading about it are terrifying. I don't want to start creating my own personal trail of use on AI. There's a lot of history of tech companies over-harvesting data. I might think I am just sharing one piece of data, but then I'm authorizing it to access all kinds of my data. I'm not willing to go there at this stage.” She added that the AI tool would only be valuable if it gave her very specific information, but that she was not willing to share her specific details, background, and location with the tool in order to get its specific guidance. The second skeptical participant (P15) had similar concerns. She explained:

I'm on the fence on it. I want it to know and understand me, my thought process. But then I don't want it to know all of that. It's like going down a rabbit hole. You want things to be easier, but then things have to understand you, and you're letting them into your personal life or mind. ... There's no middle ground. It has to record the data and store it. Who is the big brother behind the data watching it? Can they use it to affect my life in some other way?.. How might it bite you in the ass later on? (P15)

Aside from these two skeptical participants, others were generally positive about AI and its value to them. Several participants mentioned that AI is at its beginning stages, and said that they would be wary of how reliable its information would be as the platforms roll out. At the start of the survey, participants were asked to rate how much they would trust what an AI tool like ChatGPT would tell them about a life problem like their employer not paying them, their landlord harassing them, or how to figure out child custody. Participants reported that they would have medium levels of trust, with explanations like “I don't think AI is ready to give legal advice. It is too early” (P3), or “I trust it somewhat. The info is legitimate. But I know there are holes in using AI. I have used it before and gotten partial info. It's still not a complete system” (P12). On average, participant reported a 2.7 out of 6 trust rating, with a standard

deviation of 1.7.

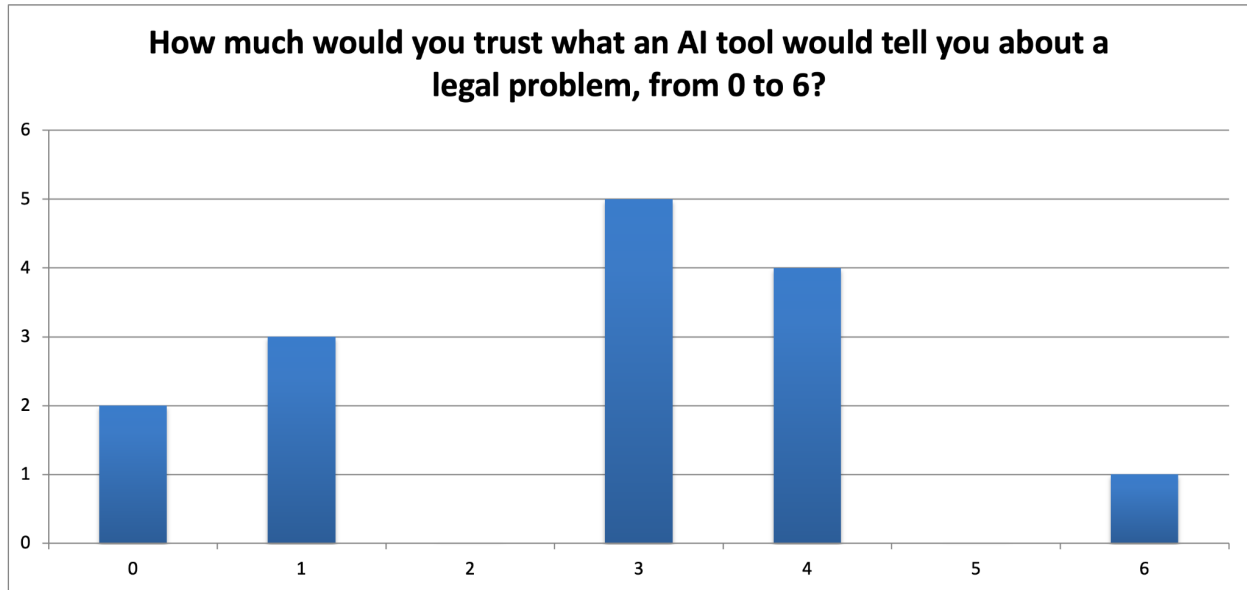


Figure 1: Reported trust of an AI tool, before using a tool

2: How Participants Interact with & Assess an AI Tool

As we observed participants use the AI tool Google Bard to deal with a fictional scenario of receiving a warning notice about an eviction, we observed how participants used the tool, what they said about it, and how they assessed the value of the tool after finishing their interaction with it.

Participants were allowed to use Bard for as much or as little as they wanted, to respond to the fictional scenario. They were instructed to pretend they were sitting in their living room, after having received the 2-week warning notice about an eviction, and use Bard as much as they would like to deal with the problem. A few participants reported that they would not use an AI tool like Bard if this scenario was really happening to them. For example, P15 said she would not use an AI tool, but rather “I’d call whoever put that notice on my door. I want to understand where it came from and what I can do about it.” P2 similarly reported that she wouldn’t turn to technology, but that if she received an eviction notice, “I would talk to building neighbors, talk to tenant union, housing rights committee, attorney colleagues.”

Most other participants, though, reported that they would go online as an initial step to deal with an eviction notice. They reported that they would usually go to a search engine like Google or directly to their city or state government’s webpage to find resources. They reported that if they were to use an AI tool like Bard, they would likely follow up the AI interactions with going to a search engine to find more information, or going to the websites and organizations the AI tool mentions to get more information and services.

Participants who had not used an AI tool before tended to enter prompts as if they were using a search engine. They entered a few words to get information from the AI tool, like:

“Tenant rights” (P15)

“Landlord issues” (P14)

“Evicted by landlord” (P10)

Five of the participants, all of whom hadn’t used AI tools, before only entered in one prompt and stopped after this interaction. Other participants, though, entered a series of short prompts, gradually getting more specific. For example, P13 went through the following six prompts while interacting with Bard:

received an eviction notice what can I do?

california state eviction law

Alameda Eviction laws

Alameda relocation assistance

Legal help for eviction notice

How do you apply for the relocation assistance for alameda county

He explained that he used the response from each prompt to then know what follow-up prompt to ask for, gradually figuring out what his options were and then asking the tool to give more information about these possible actions. This behavior was the exception, though. Most novice AI-users tended to enter a general prompt that phrased their situation into a few words, did not include their issue, and did not ask the AI tool for anything specific.

Participants who had used AI tools like ChatGPT and Bard before tended to enter more detailed prompts, that asked the tool a specific question.

I've an official legal document from my landlord that I have been evicted. What legal recourse do I have? (P12)

I would like to get in touch with a legal aid organization to assist me with what I believe is an illegal eviction (P12)

Can my landlord evict me in 2 weeks california? (P5)

Where do I go to file against an eviction notice in Maryland (P9)

Experienced participants tended to use multiple prompts. Even in their first prompt, they explained that it was important to put their location and details about their scenario. Then in subsequent prompts, they would ask about more details and follow-up from the tool’s answers. The experienced participants were comfortable rephrasing their prompt if the first answer didn’t produce information they found specific or actionable enough.

Participants did not explore the interface and options that Bard offers, aside from the question-and-answer parts of the tool. They did not try out features that would let them refresh the tool’s answer to see if it would answer differently. They did not try out the “Google It” button that would let them see related

Google searches and explore those results. Participants also did not engage with sidebars or menus. Their sole focus was on the box where they could type a prompt and the response that the tool generated.

After using the AI tool to respond to the legal problem scenario, participants were asked assessment questions about the tool's value, and they were also asked to propose changes that could improve the tool's usefulness. Generally, participants found the tool to be helpful. When asked to rate the helpfulness of the tool on a scale of 0-6, the average rating was 3.6 with a standard deviation of 2.1.




Figure 2: Post-scenario rating of the AI tool's helpfulness, on a scale of 0-6

They explained that they rated the tool as helpful because it seemed to quickly understand the scenario and intent that they were prompting it with; it provided clear, straightforward answers to the prompt; and it seemed to be providing accurate, detailed information that the participant could use to take action to address the problem.

It's pretty easy and straightforward. I didn't know it could do stuff like this. It's crazy, compared to where things used to be. It was extremely helpful. It gave me exactly the answers I was looking for. It understood it from the get-go. (P13)

The results are more simple. Google [Search] can be very overwhelming with results. Compared to what I'm used to, it's a lot more digestible and easier to read. (P5)

It was extremely helpful because it gave me a better understanding prior to asking the question. Prior to this, I didn't know anything about the topic, but now I know -- 2 weeks could be possible, my landlord might have a right, and I need to search further to know what I should do for this eviction. (P14)

Bard's useful because it outlined the criteria about whether a notice is valid. The landlord didn't give me a reason for the eviction, so it's not valid. Also, the number of days the landlord gave doesn't seem correct. I'll use this to write a letter to the landlord. (P10)

I like the way it looks, it's very easy to read, it's very cut and dry. Very direct. It's good advice, because it takes a long time to figure out about eviction. Legal aid sites come up. This is awesome. This is the stuff that comes up when I Google, what I've seen before. But this is awesome that it's all right here on one page. And the way that it's not cluttered, Google has this, this this, it's so overwhelming and cluttered. This is all that you need in one stop. This would help a lot. It was awesome. I loved the way it looked, I loved the direct answers. (P6)

Some participants, especially those who were new to using AI tools and had only put in 2-word phrases into the AI tool's prompt, complained that the tool was not as helpful as it could have been, because it gave too generic content that lacked actionable, local detail.

I'd say 3 out of 6 in terms of helpfulness. Because it wasn't really giving anything really useful. It was very basic, and I still have to do searching through that, to go to Google. It's just a starting point. It's not getting me where I need to go to the finish line. The Google Search has triggers to help you think what questions to ask, and find the answer. This just gives you a basic Wikipedia like page, with no clickable links. (P15)

The tool was somewhat helpful. It gave me some insight to the fact there are forms to fill out, agencies to explore, it gave me a starting point. But I still have a lot of questions. (P12)

Participants were also asked about how much they trusted what Bard told them, on a scale of 0-6. The average trust rating was 4.2, with a standard deviation of 1.8. This was a substantial increase from the prospective level of trust that participants reported that they would have of an AI tool, before going through the scenario exercise. Before using Bard, they had an average trust rating of 2.7, and it went up to 4.2 after using the tool.

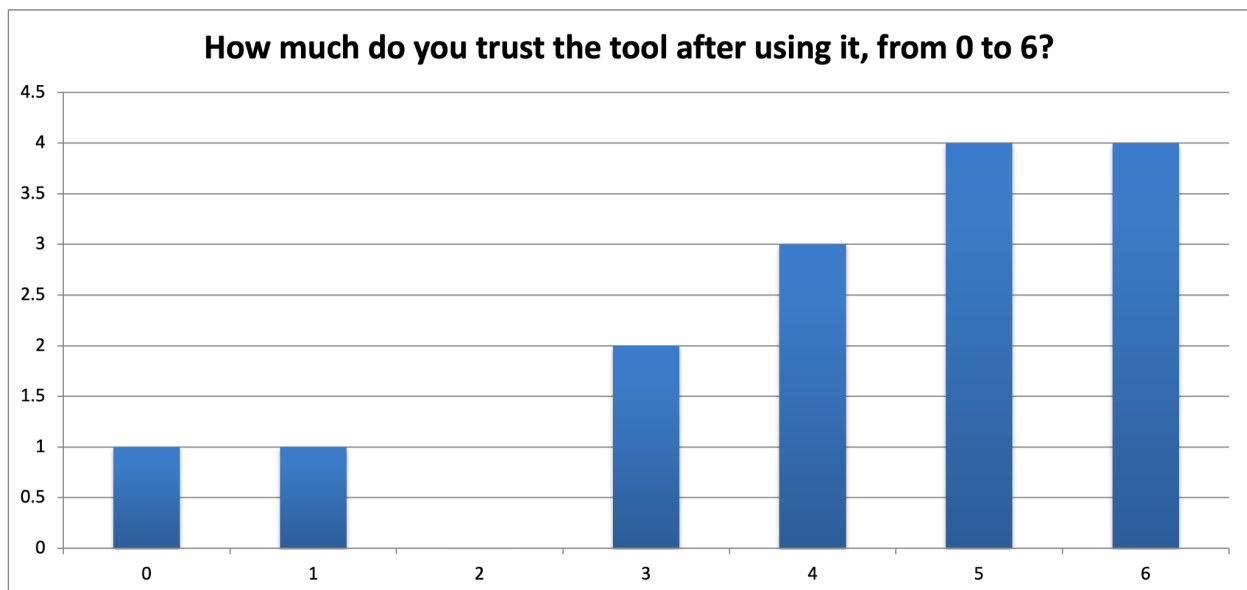


Figure 3: Trust of Bard, after using it, on a scale of 0-6

When asked to explain their trust rating after using the AI tool to address the legal problem, many participants discussed how the presentation and specificity of the tool made it seem trustworthy, as did mentions of government or legal organizations that seemed familiar to the participant.

I completely trust it. It gave me factual information and it gave me state-run website that I'm familiar with. It didn't send me anywhere spammy or anything. (P9)

In this case, I trust the response 5 out of 6. Mostly because it is literally showing me what I already thought, my priors. Nothing about it felt wrong. It didn't get so specific and start citing case laws that I had to start questioning it. It was specific to the state I'm in. If it did have citations, it would take it to completely trust, 6 out of 6. It's completely believable and in accordance with my own knowledge. (P8)

Other people also mentioned that because the AI tool was run by Google, they trust it more. The trust that they have towards Google as a search engine was transferred over to the AI tool.

I trust it, I imagine it took the information from legal and attorney websites. I think that Google has a large amount of information. When I ask the question, I'm sure it pulls the information from the results it has. I'm assuming it's valid and credible information from either court websites or law firm websites. I've never had an experience when Google gave me inaccurate information. (P5)

Google is reliable, everything in the world is fake and you need to do due diligence but I trust that on Google I would reasonably find what I'm looking for is trustworthy. (P15)

Participants also answered questions about what made the tool valuable and helpful, what it was missing, what frustrated or concerned them, and what they would change about it. The primary complaint and frustration reported was the lack of hyperlinks, citations, and sources. Participants often commented that they felt the information was right, because it made reference to details or organizations that seemed familiar and legitimate to them.

But even with this sense that the information was right, participants still were frustrated that they did not get citations about where the information came from, and that the specific statutes, cases, organizations, forms, and legal claims that the tool mentioned did not have hyperlinks to let the participant either verify that they were real and applicable, or that would let the participant immediately go to that thing to take action.

I want links to the organizations, and to the regulations it's citing. I don't want to take Bard's word for it. I might look more into the organizations. I probably would stop here. It gave me the information I wanted, even if it didn't give me the citations. I want citations because I know from my experience that they often produce made-up or incorrect info. I wouldn't trust it unless there was a link to the government or one of these organizations. (P1)

Nearly every participant complained about the lack of hyperlinks within the answers, as well as the lack of citations or websites to find more information at the end of the tool's response. These 2 changes were the top suggestion for improvement. The other frequent request was for a feature akin to "People Also Ask," in which the tool would suggest follow-up prompts that the person could ask or engage the person

in conversation to teach them how to ask more specific, effective prompts. Several people, especially those new to generative AI tools, remarked that they felt like the tool should be able to get them helpful, specific responses, but that they didn't know what prompts to enter that could get them this information they desired.

Participants were also asked about their preferences regarding simplicity versus complexity in what the tool presented them; if they would want more detailed explanation about how the tool functions; and if they would want a stronger message about the dangers of using the tool for legal questions. Several participants reported that they wanted the AI tool to offer very simple responses, at about the length and structure of Bard's current responses, but with the option to explore more detail made clear. When asked how much explanation they would want from the AI tool about why it's giving them the answers it does, on a scale of 0 to 6 the average rating was 2.9 with a standard deviation of 1.7. The value of the tool was its synthesis of many different sources and kinds of information into one executive summary.

I don't want to read a whole novel. I want the Cliff's Notes. But if it's really interesting, or I want to know more, I can get the detailed version. (P15)

I want in between simple and complex. I don't want to be overwhelmed. Give me more info like about what specific forms to fill out, not just to go fill out a form. (P12)

Keep it as it is, then people can go to blogs and more complex things down the line. Give it straight on one page, don't mess up that first page, how direct it is. (P6)

I like simple. That's the whole point of this. Otherwise I could do a Google Search myself for extensive results and information. (P5)

Several participants identified themselves as being "more is more" in regards to information when trying to solve problems like a legal one. They said they would have multiple tabs open at once, and read through as many experts' articles and people's shared experiences on social media to find their own strategy and begin taking action. Even these participants said that the AI tool should remain similar to its current interface design and complexity level for most people, but that it should have extra options for power users to dig into more legal context and also technical context.

Show me more, show me why you did this, show me your citations, show me your confidence rating (99.8 percent, 35 percent). I'm familiar with statistical analysis. I'm always of the opinion that more is more, give me all of it. The legal citations that should be in the main box, the top-level, for everyone's answer. Then the more technical jargon, the stats, the where it pulls info from, should be an additional option. That could be a subscription, Pro option to get all of that stuff. (P8)

Aside from this participant, most others said they would not want to know technical details about the AI tool's operations. Rather, the complexity that they would appreciate would be about more specific steps that the participant could take, or the citations that would verify the information that the tool was providing. As one participant said, "Just give me the answer. I don't want to hear your explanations behind it. But I do want to know where you got it. Is it factual or true? Where did you get that from? But I don't want explanations" (P15).

When asked about the value of warnings not to use the tool without consulting a lawyer, participants had mixed to negative reactions. On a scale of 0 to 6, participants rated the value of such a warning at 2.6, with a standard deviation of 1.9. Several participants said that the warning was not needed because it should be obvious to people not to rely on AI for anything akin to legal advice. Other participants reacted negatively to the idea of a warning about not using the tool for legal advice.

Don't warn me about lawyers. Most people don't have money. They also don't have time to look for an attorney and wait for an appointment. AI is quick, it gives the answer straight away. (P10)

I don't want that. If the tool warns me about needing a lawyer, I'll have a lot more doubts. It might be useful, but then I'll have to be careful. (P14)

What kind of information are you giving me that I would have to talk to a lawyer? What? (P15)

Participants tended to rate the legal advice warning disclosures as low to medium helpfulness, either because they felt it was redundant and obvious, or because they felt it was counterproductive or would give them a sense of doubt about the AI tool that they preferred not to have to deal with.

3: What the AI Tool Presents to the Participant

The research team noted down what the AI tool Bard presented to the people in response to their prompts. In only 2 instances did the tool give a non-answer response, in which it declined to give any information back with an explanation like “I'm not able to help with that, as I'm only a language model.” For every other prompt, the tool did return a response to the participant.

These responses tended to follow a similar structure. The pattern typically was as follows:

- An empathetic first sentence, like “I'm sorry that you are dealing with this,” or “Problems with your landlords can be stressful.”
- A succinct answer or statement to the question or situation that was posed in the prompt. In some cases, this was a statement of what the law says about how long a landlord has to give a tenant with a warning notice, what the steps of an eviction are, or what actions a tenant can take if worried about an eviction. This often is presented in a bullet point or ordered list.
- After this answer, the tool often recommends talking to experts like lawyers or government agencies to find out more details, and explore exactly what the best option might be.
- The tool typically concludes with a list of organizations that could have websites to read more, or services to help the person. These are not hyperlinked, but presented as a list of names. (Note that during one of our last sessions at the end of June, Bard apparently changed its interface design and began including links and image previews of websites that it mentioned. Before this change, there were no links or images in its responses.)

When asked by a participant about what legal precedent could support a tenant who wanted to fight back against an eviction notice, the tool presented a list of cases and statutes that the participant could reference as precedent. The research team noted down these cases, but could not find any record of them during subsequent searches of case law. The statutes, however, did exist as referenced by the tool, and did relate to tenant protections against eviction.

The tool frequently did mention specifics, including telling people the specific amount of warning time that a landlord in their state should give their tenant before filing for an eviction, and the specific kinds of motions and pleadings that a tenant could file to challenge an eviction in court. The tool seemed to give these specifics if the prompt asked for things like “can landlord tell me to leave in two weeks?”, “Where do I go to file against an eviction notice in Maryland?”, or “How do you apply for the relocation assistance for alameda county?” These specific prompts, that include jurisdiction, and a request to explain whether something is legal or not, or what the legal procedure is, tended to elicit a very specific, detailed answer from the tool.

When participants used short, general prompts, like “landlord issues” or “tenant rights”, the tool responded with a general answer, that did not go into specifics about timelines, process, requirements, or laws. Our team also observed that the AI occasionally misunderstood the person’s situation when the prompt was shorter, and that the tool responded with information that may not have been relevant (even though the participant was not aware of this). For instance, some participants mentioned in their prompts that they were being evicted. The tool then understood this as if a lawsuit had been filed against the person in a formal eviction suit, and gave steps for the person to respond with motions, defenses, or other court-centered strategies. The participant did not necessarily understand, nor did the AI, that the letter on their door was a notice of a possible eviction, but not a formal legal eviction action, and so their legal options in court, or in an informal resolution process, would also be different. The tool did not ask any follow-up questions about the person’s situation or the timeline of interactions.

Our team noted down what organizations the tool recommended people contact next. The tool most often recommended people to seek help from the National Low Income Housing Coalition (NLIHC). The tool also recommended participants to go to HUD, Lawhelp.org, the Legal Services Corporation, local bar associations, and statewide housing and consumer protection agencies. The tool did not recommend or link to any commercial services. It did not include any apparent paid advertising.

Discussion of What People Want & Need with Legal Help AI

The findings from this limited sample of adults point to hypotheses for further research, policymaking, and tool development. Though the findings do not offer conclusive evidence about how people will use AI tools for legal help scenarios, they do indicate some patterns that future surveys, interviews, observations, and townhalls might explore, to determine how prevalent and impactful they are.

Hypothesis 1: People will use AI tools like Bard and ChatGPT to deal with life and legal problems with increasing frequency and prominence.

The majority of our participants indicated that they found the AI tool helpful in the legal problem scenario and that they would be likely to use it if they encountered future scenarios in which they needed assistance. Most participants found value in the tool, especially once they were able to devise prompts that elicited specific statements of rights and the law, procedural guides, and referrals to other groups that could help them.

The AI tool Bard’s current responses were in sync with the kinds of content and presentation that people found valuable. The executive summary style of content fit with the need of having a short attention span, and wanting to get a sense of control about the problem area they were in. In particular, people appreciated three kinds of information:

- **Statement about the law applied to this situation.** Frequently, participants wanted the AI tool to tell them if another person's actions (in this case, their landlord) were legal, and if they had any rights or recourse within the legal system. The participants wanted this statement to be as authoritative as possible, so that they could use it as a reference point in conversations, calls, letters, or lawsuits with the landlord, or with the court.
- **Procedural Step-by-Step.** The participants also wanted a clear guide to either the options they could take (e.g., to negotiate, to sue, to move, to report the landlord to an agency), or a guide to the exact steps to follow to file a document, raise a complaint, apply for assistance, or other procedure. Participants did not want just a menu of options, they also wanted a distilled list of specific actions to take.
- **Service Directory.** Many participants remarked that they would note down the organizations that the AI tool recommended, to visit their websites and possibly call for help. Even if many of the participants felt confident in solving problems on their own using the Internet, they appreciated having a curated directory of human services to use if the DIY approach did not work.

Participants not only appreciated these 3 kinds of content that the AI offered them, they also appreciated that it was presented in a clear, uncluttered way. This combination of specific guidance and user-friendly interface design indicated that the AI was valuable to them as a starting point to understanding their problem and what to do. AI tools may become a common "first step" when problem-solving. Instead of a search engine, people may go to AI chat to get an executive summary of what their problem is, what options they have, and what to do next. As many participants mentioned, they would go from this AI tool to then explore a search engine or a social media site like Reddit. These other platforms become follow-up tool to verify information the AI has given, and to deepen understanding and support.

Hypothesis 2: Many people will over-rely on AI tool's explanations about the law, even if there are warning disclosures or statistical details.

Traditional disclosures, presented in terms of service and warning boxes to the sides of interfaces, may not be sufficient to protect many people from possible consumer harms of relying on information about legal rights, procedures, and services. Nor may more technical warnings that give information about the tool's confidence level or probability of accuracy. Based on participants' feedback, many people will not engage with these warnings both because of the burden of reading extra information, but also because they don't want to have a cloud over the content they are receiving from the AI. Several participants indicated that these warnings would upset them, because they felt the AI was giving them clear, important information to deal with a problem, and they wouldn't want warnings to 'rain on their parade'.

This potential avoidance of warnings and technical details is concerning, because there are several quality problems that we observed with the information that the AI tool presented. Though the tool effectively communicated in plain language with clear, actionable presentations, it did not always present high quality information about the substantive law, procedural steps, or services to use. This raises the concern of legal help information that looks good at first interaction, but actually is incorrect, unhelpful, or inapplicable to the person's situation. This phenomenon could be called 'ersatz legal help'.

In particular, the AI presented three kinds of ersatz legal help: outright hallucinations of legal information; correct information presented without sufficient context ; and correct but unhelpful or inapplicable information. Outright hallucinations included presenting legal cases about tenant's rights that

could not be found upon later research; it seems the AI tool created new, but real-seeming cases to answer the participant's questions about legal precedents. More frequent were instances in which the AI presented correct information without sufficient context, like when telling participants that they had a certain amount of time before facing eviction, without informing them of possible exception scenarios that would complicate this. The AI tool also frequently referred participants to organizations that do not offer any direct services (like NLIHC, which has no resources or services for tenants, but rather focuses only on policy and data), or to services that were for the wrong jurisdiction (like sending a Californian to the Colorado state bar association).

In these situations of ersatz legal help, the participant did not recognize that there was a quality issue with the information being presented. They assumed that the information was correct and applicable, though they reported that in real life they would double-check the accuracy by searching online after using the tool. They may discover after future Google searching that the case Bard told them about does not exist. They may realize when they come to court or talk with a lawyer that they have been relying on legal rights or rules for another jurisdiction, or that do not exist at all. They may try to file a motion to vacate an eviction in court, only to realize that they cannot because their landlord has not filed an eviction lawsuit against them, they have only given them a warning notice. These situations of ersatz legal help on AI platforms seems to be different than quality issues occurring on other platforms, like search engines. Because the AI platforms deliver a response in paragraphs and lists to the user, rather than a list of options on a search engine results page, users may be more likely to rely on the AI's responses as the 'answer'. More research is needed to compare people's confidence in different types of AI results, and compared to search engine results.

Among our participants, we recognized three different personas when it came to critically assessing the quality of legal information and the risks with relying on it to take an action in response to an eviction warning notice.

Type 1: "I'm Going to Screenshot This": The first type of user perceived the AI tool's responses as evidence they could use to support a message they planned to send to another party. A minority of our small convenience sample (2 out of 15) were in this persona group. These users would employ the tool's responses, such as screenshots, copied text, or attachments, in an email, letter, or text message to their landlord, using it as authoritative information to substantiate their case. They would assert their rights, stating that they would not relinquish their apartment and provide the AI tool's output as the basis for their argument. Future researchers should explore how frequent this behavior is, because it involves concerning risks of reliance on incorrect or hallucinated legal information.

Type 2: "Tell Me The Law (and I'll Cherry-Pick From There)": The second type of user sought the AI tool as a way to get a definitive answer about the law. In our small sample, we identified this type of behavior in 4 of 15 participants. This kind of participant would scan the presented content and seize upon specific details, such as a 30-day notice requirement or exceptions for unforeseen circumstances. These users were primarily looking for confirmation of their own suspicions, and the AI tool would provide them with specific information that aligned with their expectations. Some users would then verify the answer by conducting additional research, while others would directly employ the content as a reference point when engaging with their landlord or seeking assistance. Essentially, they would extract the specific

information from the AI tool's response, either by clipping it out or paraphrasing it, and assert that it represented the answer. Some users would validate the information, while others would not.

Type 3: “Now I Know What to Research”: The third group of users regarded the information provided by the AI tool as a framework rather than a definitive answer. The majority of our participants (7 out of 15) were in this persona group. This group found it valuable and helpful in guiding their understanding of the problem at hand. They expressed sentiments such as, "Previously, I had no idea how to approach this problem, but now the tool has helped me understand what questions to ask or how to consider it." However, they did not treat the AI tool's output as a statement of the law or use it as evidence of their rights or the obligations of others. Instead, they carefully read the content and recognized that they needed to conduct further research by referring to specific authoritative sources, such as government websites or trusted organizations. They acknowledged their lack of knowledge regarding the legality of a two-week notice, for example, and realized the need to engage in discussions with their landlord, review local laws, or thoroughly examine the lease agreement. The AI tool served as an educational resource, providing them with a starting point, highlighting keywords and key questions. They would then either continue using the AI tool or explore other sources to gain a comprehensive understanding of the law and its application to their situation.

Participants differed in critical thinking regarding the sources, quality, and usefulness of the AI tool. Some demonstrated an ability to think critically about information sources by explaining that they knew the importance of looking for nuances, exceptions, and local and state laws when searching online. Some participants remarked that information available online can be unreliable or spam, and that AI tools might provide them with hallucinations that looked authoritative but were incorrect. Our research team observed that these participants brought critical thinking to their interactions with AI, but still reported its value as an orientation framework. Other participants, however, were eager to make use of the AI tool's statements about the law, and said they would use it immediately without verifying it. These participants were not as aware of the importance of jurisdiction and nuance around legal information, and also were not as aware of the potential for AI hallucinations. Future researchers may add in a screening module to identify study participants' media and technology literacy, to ascertain how practiced they are in verifying the accuracy of information they receive online and how much they know about verifying legal information.

With these patterns of over-reliance on the AI's legal information and the reluctance to engage with warnings or technical details, technology platforms, legal industry, and government policy-makers must explore other ways to protect people from possible harms that might result. By thinking about specific harms, technology companies that operate AI tools might better be able to mitigate risks and protect people from bad outcomes. For example, platforms could explore this initial list of quality problems, and devise ways to adjust their model, guardrails, and partnerships with domain experts to lessen the risk of harm.

- **Bad referrals**, in which the platform presents the user with a curated list of groups who can help them, but that contains inapplicable organizations. This can be addressed by a collaboration with domain experts in legal services and courts, to identify the local groups that have expertise and capacity for a given problem type in a jurisdiction.
- **Hallucinations of legal cases**, that are presented to people who ask for precedent, legal rules, and other past situations close to theirs. AI platforms could put up a manual guardrail against

presenting case citations. If asked for legal references, they could direct people to look up cases on their own or decline to answer.

- **Incorrect jurisdiction laws and procedures**, in which the platform offers timelines, forms, and options that are correct for another location but not the user's. The tool could be programmed to respond to a legal-seeming question with a question about where the problem is happening. The tool could then consult a jurisdiction-specific model to return a response that drew from the correct location.
- **Cherry-picking legal details**, in which the user recognizes one date or legal phrase and does not pay attention to the context and warnings the tool provides. The tool offer a high-profile link, with an image or button, to encourage the person to visit a guide, form, intake line, or FAQ page from a local legal help organization for the issue area. This could more prominently direct the person to either engage with the complexity of the law, or find someone who can help them do so.

Domain experts, technology companies, and community members can work together to further identify the specific bad outcomes that might result from people using legal help AI to understand the law, find what steps to take, and seek help. They can then collaborate on new technical, interface, and policy solutions that can mitigate these risks.

Conclusions

This initial qualitative study of people's interactions with AI for legal help identifies themes, user archetypes, and hypotheses that future research can explore. The study demonstrates that it is possible to gather empirical data from community members about how they would use AI to deal with legal problems, what they find valuable and harmful, and what preferences they have for future interfaces, technical design, and policies for these AI tools.

Though this study indicated that there is no single, universal 'member of the public' to design legal help AI tools for, it does point to some patterns and hypotheses that future interviews, surveys, workshops, and townhalls might investigate further. The study's participants indicated a general enthusiasm and optimism around how AI can help them understand the law, protect themselves, and feel a sense of control and dignity. Domain experts worry about whether this enthusiasm might come at the price of low-quality information or other harms, but this worry is not helpful without clear understandings of the possible harms, of how people are likely to interact with policy measures like warning notices or technical explanations, and of the ideal solutions that community members prefer.

Future studies and policy-making efforts should concentrate on identifying these specific preferences, behaviors, and quality problems, building off of this limited initial survey. This future work might then result in a definitive risk typology, of the actions and outcomes that technology companies, legal professional groups, and government agencies must plan around, as well as the interface and technical solutions that are most likely to engage users and mitigate these risks. Ongoing, human-centered research can ensure that the potential for AI to close the justice gap can be realized, while the public is also empowered and protected.

-
- ¹ American Academy of Arts & Sciences. AI's Implications for Equitable Access to Legal and Other Professional Services. In Cambridge, MA; 2023. Available from: <https://www.amacad.org/events/ai-implications>
- ² Self-Represented Litigation Network. Generative AI April Webinar [Internet]. 2023. Available from: <https://attendeegotowebinar.com/recording/3287712821124498779>
- ³ Brescia RH, McCarthy WA, McDonald AM, Potts KB, Rivais C. Embracing Disruption: How Technological Change in the Delivery of Legal Services Can Improve Access to Justice. *Albany Law Rev* [Internet]. 2015;78(6):553–621. Available from: <http://papers.ssrn.com/abstract=2515009>
- ⁴ Telang A. The Promise and Peril of AI Legal Services to Equalize Justice. *Harv J Law Technol* [Internet]. 2023; Available from: <https://jolt.law.harvard.edu/digest/the-promise-and-peril-of-ai-legal-services-to-equalize-justice>
- ⁵ Simshaw D. Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services. *Yale J Law Technol*. 2022;24:150–226.
- ⁶ Cyphert AB. A Human Being Wrote This Law Review Article: GPT-3 and the Practice of Law. *Univ Calif Davis Law Rev* [Internet]. 2021;55:401–43. Available from: https://lawreview.law.ucdavis.edu/issues/55/1/articles/files/55-1_Cyphert.pdf
- ⁷ Kanu H. Artificial intelligence poised to hinder, not help, access to justice. Reuters. 2023 Apr; Available from: <https://www.reuters.com/legal/transactional/artificial-intelligence-poised-hinder-not-help-access-justice-2023-04-25/>
- ⁸ Stepka M. Law Bots: How AI Is Reshaping the Legal Profession. *ABA Business Law Today* [Internet]. 2022 Feb; Available from: <https://businesslawtoday.org/2022/02/how-ai-is-reshaping-legal-profession/>
- ⁹ Blomkamp E. The Promise of Co-Design for Public Policy. *Aust J Public Adm*. 2018 Dec 1;77(4):729–43.
- ¹⁰ Hagan M. Participatory design for innovation in access to justice. *Daedalus* [Internet]. 2019 [cited 2019 Apr 24];148(1):120–7. Available from: https://www.mitpressjournals.org/doi/pdf/10.1162/daed_a_00544
- ¹¹ Hagan MD. A Human-Centered Design Approach to Access to Justice Generating New Prototypes and Hypotheses for Intervention to Make Courts User-Friendly. *Indiana J Law Soc Equal* [Internet]. 2018;6(2):199–239. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3186101
- ¹² Organisation for Economic Cooperation and Development. Access to Justice and Legal Needs Surveys. In: *Legal Needs Surveys and Access to Justice* [Internet]. Paris: OECD Publishing; 2019. Available from: <https://www.oecd-ilibrary.org/docserver/cab05cff-en.pdf?expires=1687203933&id=id&accname=guest&checksum=0F2AEFC895E941489CC5615B6FCB283D>
- ¹³ Sandefur RL. What We Know and Need to Know About the Legal Needs of the Public. *S C Law Rev*. 2016;67:443–59.
- ¹⁴ Legal Service Corporation. Justice Gap Study 2022 [Internet]. Washington, DC; 2022 [cited 2022 Oct 12]. Available from: <https://justicegap.lsc.gov/resource/executive-summary/>
- ¹⁵ Sandefur RL. Accessing Justice in the Contemporary USA: Findings from the Community Needs and Services Study [Internet]. SSRN. 2014 [cited 2019 May 23]. Available from: http://www.abajournal.com/files/sandefur_accessing_justice_in_the_contemporary_usa_aug2014.pdf
- ¹⁶ Hiil, IAALS. Justice Needs and Satisfaction in the United States of America [Internet]. Denver, Colorado; 2021 [cited 2021 Sep 1]. Available from: <https://iaals.du.edu/sites/default/files/documents/publications/justice-needs-and-satisfaction-us.pdf>
- ¹⁷ Sandefur R. The Importance of Doing Nothing: Everyday Problems and Responses of Inaction. In: Pleasence P, Buck A, Balmer N, editors. *Transforming Lives: Law and Social Process* [Internet]. Stationery Office Books; 2007. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1599755
- ¹⁸ Cabral JE, Chavan A, Clarke TM, Greacen J, Hough BR, Rexer L, et al. Using Technology to Enhance Access to Justice. *Harv J Law Technol*. 2012;26(1).
- ¹⁹ Staudt RW. All the Wild Possibilities: Technology That Attacks Barriers to Access to Justice. *Loy LAL Rev* [Internet]. 2008;42(10). Available from: http://heinonlinebackup.com/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/lla42§ion=43

-
- ²⁰ The Engine Room. A Global Review Technology for Legal Empowerment [Internet]. 2019 [cited 2019 Jul 3]. Available from: <https://theengineroom.org>.
- ²¹ Greacen JM. Eighteen Ways Courts Should use Technology to Better Serve Their Customers. *Fam Court Rev* [Internet]. 2019 [cited 2018 Dec 19];57(4):515–38. Available from: http://iaals.du.edu/sites/default/files/documents/publications/eighteen_ways_courts_should_use_technology.pdf
- ²² Sunday K. There’s potential for AI chatbots to increase access to justice. Thomson Reuters Forum: AI & Future Technologies [Internet]. 2023 May; Available from: <https://www.thomsonreuters.com/en-us/posts/legal/forum-spring-2023-ai-chatbots/>
- ²³ Holt AT. Legal AI-d to Your Service: Making Access to Justice a Reality. *Vanderbilt J Entertain Technol Law* [Internet]. 2023; Available from: <https://www.vanderbilt.edu/jetlaw/2023/02/04/legal-ai-d-to-your-service-making-access-to-justice-a-reality/>
- ²⁴ Wu J. AI Goes to Court: The Growing Landscape of AI for Access to Justice. *Legal Design and Innovation* [Internet]. 2019 Aug; Available from: <https://medium.com/legal-design-and-innovation/ai-goes-to-court-the-growing-landscape-of-ai-for-access-to-justice-3f58aca4306f>
- ²⁵ Tito J. How AI can improve access to justice [Internet]. 2017. Available from: <https://www.centreforpublicimpact.org/insights/joel-tito-ai-justice>
- ²⁶ Rostain T. Techno-optimism & access to the legal system. *Daedalus* [Internet]. 2019 [cited 2019 Apr 24];148(1):93–7. Available from: https://www.mitpressjournals.org/doi/pdf/10.1162/daed_a_00540
- ²⁷ Guzman H. AI’s “Hallucinations” Add to Risks of Widespread Adoption [Internet]. *ALM Law.Com: Tech Sector Risk Watch*. 2023 [cited 2023 Jun 19]. Available from: <https://www.law.com/corpcounsel/2023/03/23/ais-hallucinations-add-to-risks-of-widespread-adoption/?slreturn=20230519164801>
- ²⁸ Buchholz RM, Chivers CE, Fiedler ND, Johnson AA, Schnake KG, Storey JL, et al. Regulation of the legal profession in the United States: overview [Internet]. 2021. Available from: [https://uk.practicallaw.thomsonreuters.com/2-633-6340?transitionType=Default&contextData=\(sc.Default\)&firstPage=true](https://uk.practicallaw.thomsonreuters.com/2-633-6340?transitionType=Default&contextData=(sc.Default)&firstPage=true)
- ²⁹ Harden S. How Would Regulating Generative AI Even Work? Team Do Something [Internet]. 2023 Apr; Available from: <https://samharden.substack.com/p/how-would-regulating-generative-ai?>
- ³⁰ Tripp A, Chavan A, Pyle J. Legal Services Community Principles and Guidelines for Due Process and Ethics in the Age of AI [Internet]. 2018. Available from: <https://docs.google.com/document/d/1voL045GVM9HJ2hTFq76l8aEvHOJBX8OPjwOjgwttDUU/edit>
- ³¹ Google. Generative AI Additional Terms of Service [Internet]. Privacy & Terms. 2023 [cited 2023 Jun 19]. Available from: <https://policies.google.com/terms/generative-ai>
- ³² Granat R. ChatGTP, Access to Justice, and UPL [Internet]. *The Law Product Makers*. 2023 [cited 2023 Jun 19]. Available from: <https://www.lawproductmakers.com/2023/03/chatgtp-access-to-justice-and-upl/>
- ³³ Blomkamp E. Systemic design practice for participatory policymaking. *Policy Des Pract*. 2022;5(1):12–31.
- ³⁴ Bason C, Botero A, Saad-Sulonen J, Pillan M, Suteu I, Staszowski E, et al. Public and Collaborative: Exploring the Intersection of Design, Social Innovation, and Public Policy [Internet]. Manzini E, Staszowski E, editors. New York, New York, USA; 2013 [cited 2018 Jan 3]. Available from: http://nyc.pubcollab.org/files/DESIIS_PandC_Book.pdf
- ³⁵ Dantec CL, DiSalvo C. Infrastructuring and the Formation of Publics in Participatory Design. *Soc Stud Sci* [Internet]. 2013;43(2):241–64. Available from: <http://sss.sagepub.com/content/43/2/241.abstract?etoc>
- ³⁶ Groves L. Exploring the role of public participation in commercial AI labs [Internet]. 2023. Available from: <https://www.adalovelaceinstitute.org/project/public-participation-commercial-ai/>
- ³⁷ Berditchevskaia A, Malliaraki E, Peach K. Participatory AI for humanitarian innovation: a briefing paper [Internet]. London, UK; 2021. Available from: <https://www.nesta.org.uk/report/participatory-ai-humanitarian-innovation-briefing-paper/>
- ³⁸ Yang J, Park T. The Limits of Participatory Design in AI [Internet]. *Partnership on AI*. 2020 [cited 2023 Jun 13]. Available from: <https://partnershiponai.org/methodsforinclusion/>
- ³⁹ Hossain S, Ahmed SI. Towards a New Participatory Approach for Designing Artificial Intelligence and Data-Driven Technologies. In: *CHI ’21* [Internet]. 2021. Available from: <http://arxiv.org/abs/2104.04072>

-
- ⁴⁰ Moss MA. The Escambia Project: An Experiment in Community-Led Legal Design. *Des Issues*. 2020 Jun;36(3):45–60.
- ⁴¹ Cornett L, Knowlton NA, Swearingen J, Houlberg M. Redesigning Divorce: User-Driven Design for a Better Process [Internet]. Denver, CO; 2019 [cited 2020 Jul 7]. Available from: <https://iaals.du.edu/publications/listen-learn-lead>
- ⁴² Hagan M. Legal Design as a Thing: A Theory of Change and a Set of Methods to Craft a Human-Centered Legal System. *Des Issues* [Internet]. 2020 Jun 24 [cited 2020 Jul 10];36(3):3–15. Available from: https://doi.org/10.1162/desi_a_00600
- ⁴³ Zimmerman J, Forlizzi J, Evenson S. Research through design as a method for interaction design research in HCI. *Proc SIGCHI Conf Hum factors Comput Syst - CHI '07* [Internet]. 2007;493–502. Available from: <http://portal.acm.org/citation.cfm?doid=1240624.1240704>
- ⁴⁴ Denvir C. What is the Net Worth? Young People, Civil Justice and the Internet [Internet]. University College London; 2014 [cited 2019 Nov 7]. Available from: <https://pdfs.semanticscholar.org/b584/c82bbc1baebd435a36ac1aa25001930344fa.pdf>
- ⁴⁵ Hagan M. The User Experience of the Internet as a Legal Help Service: Defining standards for the next generation of user-friendly online legal services. *Va JL Tech* [Internet]. 2016;20(2):395–465. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2942478