How to Use Artificial Intelligence in Resettlement Work: Opportunities and Challenges

September 17th, 2026

Transcript

Introduction

Krysti Nellermoe: Well, welcome to our webinar today. My name is Krysti Nellermoe. I'm the Senior Program Officer for Emerging Technologies and Digital Literacy on the Switchboard team. I have been with IRC for 11 years, and Switchboard for 3, leading all the innovation and digital literacy technical and training assistance. And I will hand it over to my co-presenter today.

Kimberly Inez Mays: Hi there, I'm Kimberly Inez Mays. I am the Senior Program Officer for Digital Platforms and AI. I oversee and manage all of Switchboard's digital properties, and I have 20 years of technology experience specializing in AI adoption and solutions development. So, most of what I do now is teach the field how to use these tools.

KN: We want to welcome you all today to this session, whether you are starting from scratch and looking at where you can incorporate AI into your practice and your work, or you've been experimenting with AI and wondering where it's responsible and where it's an appropriate use. Everyone is welcome. We are going to be hitting the whole gamut of that spectrum, and we're hoping to give you a practical, grounded foundation built around tasks you already do, like drafting case notes or communicating with clients across language barriers. Kim is going to give us some demonstrations on how to use AI in tools commonly used in resettlement work, and we will close with a three-step framework you can use to decide whether a specific task is a good fit for AI. The through-line for everything in this presentation is your expertise, your judgment, and most importantly your relationship to the people you serve — that is not replicable by AI, and today we're looking at how these tools can fit, and where they don't. So we're very excited to have you all part of this session. Our learning objectives for today are: first, to be able to describe generative AI, its capabilities, and its limitations; second, to identify two resettlement work tasks that AI can support; and third, to apply a three-step framework to determine AI suitability for daily resettlement tasks. Let's dive into Section 1.

1. Describe generative AI, its capabilities, and its limitations

KN: So, describing generative AI, its capabilities, and its limitations. But before we jump into what those are, let's get a little pulse check from you all on where your starting point is. If you can grab your phones, use the camera feature on your phone to take a picture of that QR code, and it should pull up a Slido survey for you. You can also type in slido.com with the number on the screen. And just give Kim and I a sense of where we're

KM: starting.

KN: today. How often do you use AI in your work? We've got a never, I love it. Sometimes, often, seldom — just give us that baseline of where you think you're starting today.

KM: Yeah.

KN: It's a balanced group.

KM: Yeah?

KN: Alright, never's going slower. Lower, lower, lower — up, got another person. Okay, this is a really good baseline. So, a majority of you are sometimes and often. Seldom and never are not too far behind. So, I think this is the right level for today's webinar — this is super helpful. So, Kim, let's hand it over to you, and let's dive into some vocabulary. What are we talking about here when we talk about AI, and generative AI specifically?

KM: Okay, so let's talk about the basics of AI — just some quick terminology you need to know. First, AI generates content, makes predictions, and supports decisions. These are systems that perform tasks that normally require human intelligence, like reasoning, learning, and decision-making, and they're shaped by the people and systems that train them. Generative AI, or GenAI, is a subset of AI that creates new content, such as images, text, video, or code — think ChatGPT, or Claude, or whatever you're using now; most likely it's a GenAI tool. And a large language model, or LLM, is an AI model trained on massive amounts of text and other material to understand and generate human-like language. LLMs power many GenAI tools, like chatbots. Many modern LLMs also work through a problem step-by-step before answering — these are called reasoning models, and they're slower, but better for complex, multi-step tasks.

Not all AI tasks carry the same risk. You have lower-risk tasks and higher-risk tasks, and discernment is knowing which methods apply for each level of risk. Lower-risk tasks still require human oversight, but less so than higher-risk tasks — these include summarizing a long document, turning rough notes into a first draft you later vet and refine yourself, drafting an email or message, brainstorming options or explaining a concept, and outlining or organizing your own work. Higher-risk tasks require deeper oversight, sometimes from other stakeholders — these include drafting content for reports or proposals, working with facts that others will rely on, citing research as evidence, working with local or cultural context, translating content for clients, and finalizing anything that's shared. AI can sound confidently wrong, so vet the outputs appropriately — just because it seems well-written and confident doesn't mean it's correct, and you're responsible at the end of the day.

Here are some of the limitations of AI. It's trained on images, audio, text, and more, so it reflects the choices and biases of the humans and data behind it. If a large language model is trained mostly on Global North or Western cultural contexts, you have to make sure it's not biased, and that you're adjusting and prompting for that. It also has a knowledge cutoff, so it only knows facts up to a certain point in time — that's more true of older models, though newer models can also access the internet, which helps bridge that gap. Speaking of internet access, check whether the tool is actually searching the internet for you, as opposed to relying only on its own training.

With all that said, you want to use AI responsibly. First, follow your organization's AI policy — no matter what Krysti and I tell you, what matters most is your own organization's policy, and if you don't know whether your organization has one, ask. Ask before using any tool for work tasks. Never input a client's personally identifiable information, or PII — things like full names, case numbers, dates of birth, or immigration status — into an AI tool. A lot of what you put into these tools can be repurposed to train the model, so that data isn't private; treat it as confidential and not shareable. And lastly, clients have a right to know if AI is used in their services, so get informed consent.

KN: And at the end of this webinar, we'll be sending out some resources. We recently published a blog on talking with clients about consent for AI use, and we also have a toolkit on AI readiness for organizations that steps you through determining your readiness. Please check out those resources at the end — they go more in depth than this webinar can. Something I really want to touch on, Kim, is using AI more sustainably. Here are a few thoughts on how to do that. First, match the model to the task: inside your chat feature, whether it's Claude, Copilot, or whatever you're using, you can toggle through models, and every model uses a different level of energy. Lighter models use less energy and can be great for lighter tasks, like rewriting plain language, changing the voice of a document you're uploading, or a simple translation — adapting your needs to the model is an easy way to save energy. Second, be deliberate with each query: if you've given context earlier in the same chat, the AI will remember it, so you don't have to repeat it every time — that saves energy and cuts down on back-and-forth. Do it once, or a few times, rather than many slower iterations, so each query counts. And third, reach for AI when it adds value — not everything needs AI, and being aware of the tasks that do and don't is part of using it appropriately.

And with that, let's jump into another Slido question. Get your phones out — you can use the QR code with your camera, or type in slido.com. Think about an AI tool you've used: what is one thing it does well, and one thing you'd want a human to check?
KM: Yes.
KN: It's a ride.
KM: But I'm seeing an overall theme here — it's really just creating new content, summarizing content, and doing research. I think we all pretty much use AI day-to-day in the same way, to make our lives and work easier.
KN: Yeah, humans must check sources — I think that's a great point. Don't just have something generated and ship it; take the time to go through it, check the links, check the facts. Like Kim said earlier, sycophancy is a real thing — it's going to make you sound wonderful, praise you, and sound super confident, but it always takes that extra step to refine and check. Yeah, I'm seeing a lot of that. Emails.
KM: Sorry.
KN: Beautiful — looks like you all have a pretty good sense of the things you'd want AI to do, and where you'd want to bring in that human response. There's just a—
KM: Do we summarize notes to draft a summary? And also, I see an interesting point that says it 'makes you weak in thinking.' We want to touch back on that too — do you need to use AI for this task, versus just doing it yourself? So we're seeing a lot of human in the loop, a lot of human brainwork. I'm really glad that point was made.
KN: Great. Let me just read this next. Let's move to Section 2.

2. Identify two resettlement work tasks that AI tools can support

KN: Section 2 — we're going to identify two resettlement work tasks that AI tools can support. Keep those phones out, let's get back into Slido: you can use your phone with the QR code, or type in Slido.com — what work tasks are you currently using AI for? You all mentioned an amazing list of ways you'd want humans to review and be part of that collaboration with AI, and we've got: none — a few nones, drafting policies, grant writing, language translation, developing policy, email reports, drafting policy, suggestions for client outreach, fine-tuning emails, intakes, language translations, idea brainstorming, and lots of drafting. We're seeing plain language writing, drafting SOPs — that's a great use of that. Translation. I draft documents, and then feedback I receive about the documents, and then ask to make it improved; feedback on it, yeah. Ensuring alignment, synthesizing work history. Office signs. And I want to acknowledge the folks who put 'none' — I think that's just as valid and important to note; right now maybe there aren't uses for AI in your work, and that's okay too. Refining, having fun with staff creating funny content—
KM: It's a good use case.
KN: Really good use case, yes. Drafting resumes — oh yeah. Great, I'll let just a few more come in. Event reports. Great. Well, let's dive back into some key terms, as we're getting the same language and vocabulary around AI.
KM: Okay, I'm going to go through some key terms — I'll go a little slower, since it can get technical, but don't worry if it doesn't all make sense right away. Most of you said you use AI for document creation or research. What you put into a generative AI tool is called a prompt — these are the instructions you give it: what are you asking it to do, and have you given it clear instructions to produce a useful output? As I like to say — garbage in, garbage out, though I don't love that phrase. Think of it like talking to a junior employee: if you give it a complex task with insufficient instructions, you won't get a good output. It's the same with AI — you have to make sure you understand the context you're giving it. Context is the background information: what does the tool need to know to give you a good output? Who is the output for? What tone and format should it have — a spreadsheet, a table, a bulleted list, explanations? You can be this specific in your prompt, and context can also include files, folders, and other applications your tool is connected to, including email or cloud storage. So context isn't only what's in the prompt — it's everything that accompanies it. And context engineering is the

practice of giving the AI model the right context, in the right order, for the task — being specific about the result you want and how it should use the context, and leaving out anything irrelevant. What role should the AI play, and what does it need to know to be useful?

Okay, so now we have an output — how do we work with AI outputs? I mentioned this term before: human in the loop. This means a human — you, a colleague, or anyone else you work with — reviews and approves generated content before it's used. If you're writing a draft, somebody needs to review it before it goes out to a funder, senior leadership, or another stakeholder. Any output, including spreadsheets or files derived from these tools, should be vetted. This operates at the level of an individual task or output — if I'm putting together a project plan and the AI gives me an Excel file, I need to make sure everything in that file is correct before sharing it with the project team.

We've talked about AI being confidently wrong, so let's talk about hallucination. A hallucination is sneaky, because the tool won't tell you it's making something up — it will give you facts and figures confidently, without telling you how it arrived at that conclusion, though sometimes a reasoning model lets you trace it. Every now and then it will stitch things together into an answer that may or may not be accurate, so you want to vet the outputs and watch for the model drifting or going on tangents. If that happens, course-correct: ask how it arrived at the answer, and if it feels wrong, it probably is. And then there's sycophancy — when AI praises you. These tools are designed to keep the user engaged and happy, so unless you specifically tell it to push back, it will generally go along with your ideas. That's fine if you're rewriting an email or summarizing a document, but if you're brainstorming and want to know your blind spots or biases, you have to check for sycophancy — use specific language in your prompt to tell it to challenge your ideas, and make sure it isn't just agreeing with you.

Alright, let's talk about tasks AI can support — starting with case notes. You're working with a client, and you capture a voice note or voice-to-text right after the session — with consent. AI can transcribe the recording, pull key details, and flag anything that's missing. I'm going to show a quick demo of what this looks like, and then, as a human, you review, correct errors, fill in the gaps, and finalize the case notes before they go into whatever software you're using.

KN: Yeah, let's look at this in a case study, Kim. Amar is a caseworker at a resettlement agency who has a lot of clients coming in. He just completed a 45-minute intake appointment and recorded a short voice note right after the session, while the details were fresh, because he needs to write a case note into the agency's database before his next client arrives in 20 minutes. Amar uses an AI tool to transcribe his voice note, which does not contain any client PII, and draft an initial case note. The tool flags two incomplete details. Amar reviews the draft, corrects the two factual errors, fills in the gaps, and removes one detail he didn't intend to include — he didn't just record and send, he did that second step of human in the loop. I'm going to hand it over to Kim to show this in a live demo — and I also want to encourage folks to add questions to the chat. At the end of our content, we'll be available to answer them, whether they came up during the vocabulary, the limitations and capabilities, sustainability, or the use cases in resettlement. With that, Kim, let's see how we'd go about this client case note using a voice note.

KM: Okay, so what you're seeing on screen now is Claude, but remember — use whatever tool is approved by your organization. Let's say Amar just finished his client session: he's uploading the transcript into his approved tool — in this case, for the demo, it's Claude, but use whatever your organization approves. He asks it to produce case notes, and you can see it generating something; the prompt can also specify how the case notes should be formatted. If Amar is reading the notes and wants to correct something, he absolutely can — he can say, 'actually, this should be...' and the tool will make the fix.

Let's talk about another task AI can support — I saw this come up in the Slido results too — grant proposals. You upload the RFP into your program notes, and the AI lays out the requirements and fills in the gaps in your draft, or provides a checklist of what's needed, so you make sure everything is covered and correct before sending it to the next step. You review, confirm accuracy, and strengthen the weak spots — same pattern: AI drafts, human reviews, no matter what.

KN: And before, I said 'chat' but I meant the Q&A field — please add questions there; our chat feature isn't actually active right now. Alright, let's take that proposal scenario and look at it in action. Meet Sara: she directs a small community-based organization serving East African newcomers and is applying for a foundation grant. English is her third language, and she's never written a grant proposal before. She uploads the funder's request for proposals (RFP) and asks an AI tool to summarize the key questions, required sections, and word limits. With that outline in hand, Sara writes a rough draft in her own words, then uploads it and asks the AI to improve the clarity and structure of the language and flag any errors. She reviews the suggestions, makes corrections, adds a program example only she would know, and shares the revised proposal with her supervisor before submitting.

Now, some key things to think about, because proposals are an important human-in-the-loop scenario. To give the AI context for the draft, Sara also uploads potential success stories or previous proposals she's used in the past, so the organization's voice comes through and the AI has context from past work — this improves the authenticity of the output. She identifies where a success story or partnership has been a great example for the proposal, bringing in context only she and her team would know, and that's what helps it stand apart. The tool is really just a collaborator on quality, helping make sure the draft hits everything the RFP is asking for and the word count. She is the expert, and she brings her expertise to the revisions — as AI gives her feedback or thoughts, it's her who brings those back into the draft, refining it and making it hers. When you're working with AI, there's a lot of back and forth — think of it as a collaborator, not a complete delegation of the task. You can ask: where am I missing something? Where can you push back on my thinking? Where can I improve this sentence, or the way I connect my needs, outputs, and impact stories? Over time, the tool starts to understand the organization's needs and funding priorities, and becomes a better tool the more you use it.

One really important thing: the final result is Sara's, not AI's. When she submits it to her supervisor, it's her name on it, and she has to stand behind what's written — just generating a first draft and submitting it wouldn't really be hers. Whatever you submit, you have to stand behind it, with that human-in-the-loop review, judgment, and discernment. And with that, let's see this in action, Kim.

KM: Okay, I'm doing another voiceover — again, this is Claude for the demonstration, but use whatever your organization has approved. I'll talk a bit about how the data is trained if time allows, but right now Sara is putting in the RFP details she needs to address in her draft. Because this is a demo, my prompt is shorter than it would be in practice, but you get the idea: she asks it to summarize the RFP into a checklist of required actions and word limits, and it gives her a checklist back. It depends on what you want from the prompt, but you want to make sure you've captured everything. There will be some variance in the AI's output, because it's generative, not deterministic — you'll get a different answer each time, but if you structure your prompts well, you'll get an output that serves you. In this case, she's asking, 'what do I need to include?' — and if she uploads a draft, she can also ask what's missing from it. So she's not generating new content here; she's using it as a thought partner, a check. She can say, 'challenge me on some of the content in this draft,' or 'is there anything that might be unclear here?' There are a lot of ways to structure this use case — not just telling AI to write something for you, but having it evaluate something that already exists, checking for factual issues, bias, or anything unclear that would cause a human reviewer to pause. Of course, if she takes the AI's suggestions, the product still needs to be reviewed by a human — again, human in the loop. What you're seeing here is a general, basic use case, but it can go much further with the things I've recommended.

KN: I'm seeing questions come through the Q&A feature — keep them coming, we definitely want to address them throughout the presentation. And to close out Section 2, we have another question we'd like your input on: pull out your phones and take a picture of the QR code, or type in the Slido question — after observing these two demos, what is one task in your work you'd want to try this on? Case notes — I saw some questions about case notes too, we'll talk about that toward the end, I think that's really insightful. Yep, case notes, general notes from meetings — I use that all the time.

KM: Yes.

KN: Translations — I saw a question about translation tools, we'll try to get to that. Well — yes.
KM: Grant writing, email summaries — yep, love it. Reports — oh, and a note about being very concerned with ethics — yes, we're going to talk about that.
KN: I like the design template for 101s. Keeping things consistent. CFT meeting notes, resumes.
KM: And I wanted to acknowledge a bullet point that flew by fast — there was one honest response, and thank you for it, saying they still wouldn't encourage anyone to use AI. Although we're showing AI use cases, we appreciate all opinions and points of view, because that matters to us too — so thanks for sharing.
KN: And to be mindful — if clients ask that AI not be used in their programming, exactly, that's an answer someone can give, and programming and services shouldn't shift for them because of it. It's about whether it makes sense for the task you're—
KM: Exactly.
KN: —and want to enhance with that tool. It's a tool. It's a technology. Short-term and long-term planning for team monitoring — interesting.
KM: Quick one.
KN: That's a really good one.
KM: And I really like the step-by-step instructions and guides idea, because that's what I use AI tools for. One hack I use: I record meeting transcripts, especially when I'm offboarding or transferring knowledge, and I'll use that transcript in AI to create a standard operating procedure or reference document I can use later or pass off — it saves time. It doesn't always have to be used in client work; it can be used in day-to-day administrative work too. If you're feeling uneasy about using it in client work, you absolutely don't have to, and you shouldn't if it doesn't feel right — you can use it for administrative tasks instead.
KN: Great point, Kim. I've also changed my instructions in my AI tools to push back on my ideas — where am I being too narrow, where am I creating bias? I've trained it to apply universal design, ethical frameworks, and a trauma-informed lens, so it asks me key questions to reflect on when I'm using it. It's customizable to your own needs and preferences, which makes it a pretty incredible tool, because it's not one-size-fits-all either. Creating small websites and apps for administrative work — nice!
KM: Yeah. Love it.
KN: We'll have you run the next one, whoever runs it. Alright, let's keep going into our last section.

3. Apply a three-step framework to determine AI suitability for daily resettlement tasks

KN: We are going to apply — so we're going to apply a three-step framework to determine AI's suitability for daily resettlement tasks. Keep that phone out — we have another Slido. Did I go too fast? I'm sorry, I jumped to slide 3 — we're not there yet, we're on common applications of AI in resettlement work. My apologies, my Slidos got mixed up. Over the last year, Kim and I have been fielding technical requests and training opportunities around AI and resettlement, and these six have come to the surface as some of the most common ways folks in resettlement are using AI: case notes and documentation, which just came up in that Slido; intake and triage chatbots, less common but definitely in use; affordable housing searches; translation, interpretation, and language access; program data summaries, helping with M&E processes; and partner and resource mapping — seeing what exists in a space and having AI do some of that legwork to connect clients to resources in specific communities. Now, bust out your phones again — we're going to take this list and see which common applications feel most relevant to your organization right now.
KM: Okay, so we're seeing a lot of translations, a lot of case notes and documentation, data summaries, and job search — yes, if you're searching across the internet for jobs or housing, you don't need to use data, just use a zip code and ask for jobs or housing, and you can definitely do that. Research notes, translations, resume makers — that's interesting too. Partner and resource mapping — I'm seeing a lot of cool themes and applications here.

KN: Great use case.
KM: Exactly. It really helps with new onboarding.
KN: Over and over, those tasks you see on repetition are where AI can create and automate that work, or generate the SOPs and documentation you'd otherwise have to recreate each time — it's a great way to use AI effectively. Targeted job searches — I love that. You can train AI on your company's specific style guides too, so you don't have to replicate it over and over. Now we're in the right place for applying our third and last learning objective: applying a three-step framework to determine AI suitability for daily resettlement tasks. Here's where you can take the framework and start applying it to see if a task is a good fit for AI. With that, we have another Slido: what factors do you consider when deciding whether to use AI for a task? Before we give you the framework, what frameworks do you already use? What reflective questions do you already ask yourself? 'Is it for internal use or external partners' is a great first answer — thank you to whoever put that. Is there a deadline, or not? Urgent — can I do it on my own? I love these reflective questions.
KM: Yeah? HIPAA compliance too — that's a good one.
KN: This is important, because folks on this call might not all be in resettlement specifically — clients that are new arrivals are here under 5 years under the ORR, kind of eligibility criteria. 'Is the tool I'm using HIPAA-compliant' is an important question — you should always know that. Our AI Organizational Readiness Toolkit, which Switchboard published last year, talks a lot about the levels of tools you're using: is it enterprise, through your organization? Is it a free tool? Does it cost a monthly fee? Each has different levels of security, so knowing what tool you're using, and its security posture, is on the user to know. And 'what AI tools are available to me' — that comes up a lot with organizations, making sure folks know what they can and can't use. That's a good question to bring back to leadership in your offices. Let's go through some vocab, Kim.
KM: Alright, let's talk about ethics and accountability, since we've been touching on it throughout. Ethics, as it relates to AI, means the principles guiding the responsible design and use of AI to ensure fairness, transparency, and protection of human rights, without reinforcing unfair outcomes, compromising privacy, embedding hidden bias, or creating barriers to services — it's just applying ethics through an AI lens, the same as you would anywhere else. Human oversight is the broader system of policies, review, and accountability covering how AI is used across your organization — again, if you don't know, ask your IT group or whoever manages your AI policies, to make sure you understand how you're using it safely and preventing privacy or consent issues; this operates at the level of programs and institutions. And lastly, informed consent is a client's voluntary, specific, fully informed agreement to the use of AI with their personal data — including what's collected, how it's used, who may see it, and how it's protected. Again, informed consent means a full yes, not just a maybe, especially when using sensitive data.
Okay, let's talk about the three-step framework. First: is the task a good fit for AI? Not all tasks are — some people are uncomfortable using it for case notes, and that's absolutely okay. If it's not a good fit for you, don't use it, and if the task is personally sensitive or hinges on a big decision, don't use it. Many organizations have their own criteria and policies for when AI can and can't be used, so always check there too — if it's a good use, continue; if you're unsure, don't. Second: what data am I inputting? Does the data align with your organization's content and privacy policies? It's also a gut check — if you're not okay putting it out there, don't. If yes, continue; if no, stop or adapt — anonymize the data if that better fits the policy. And third: who reviews the output? Do I have at least one human reviewer, other than myself, with relevant expertise? That matters because AI can sound confident while being wrong, so if you don't have the knowledge to check its work, bring in someone who does. Do I need another human reviewer to eliminate bias? If yes, continue with AI and human review; if no or unsure, stop and identify the right reviewer.
KN: Thanks, Kim. Here's our final Slido, wrapping up Section 3 — get your phones out. Let us know your response to this scenario: a caseworker wants to paste a client's intake notes, which include the client's PII, into a free AI tool to draft a summary. This came up in some of the questions — where does this get stopped by the framework? As a reminder, it's not on the slide, but: is this task a good fit for AI? Consider the stakes, the

complexity, and whether AI can help. What data am I putting in — does it include client-specific private information, or public information — and can it be shared? And then, who's reviewing that output? Oh, can we go back real quick? I'd love to see those responses before we jump to takeaways. Great — 68% said the data may not be secure in this tool, and that's right. The key clue in this scenario was 'free AI tool' — because it's free, the data you input is used to train the model. It's really important — the onus is on you, as the user, to know how the data is used and where the privacy policies land for a free, paid, or enterprise-level tool. If you don't know, I'd push that up to leadership to make sure it's known in the office, since AI tools have varying privacy levels.

KM: Yep, even the paid ones — you have to check.

KN: Alright, thanks — great job on that. 74% is a pretty large number, you all got the framework. Let's land the plane on our AI webinar, and we'll try to address some of these great questions in our Q&A. Next slide, please. So, our key takeaways from today: we're hoping that whether you use AI tools or don't, you're walking away knowing where AI can be impactful for your work; that you run the framework every time — is this appropriate, does it use private data, do I have a human in review; and that AI is a drafting partner, not a decision-maker — it's always a collaborator, and you're always in the driver's seat. And remember, comfort with AI starts with curiosity: ethics is what keeps it safe, and knowledge and skill are what turn it into an incredible tool for your work. Any technology can feel a little overwhelming at first, but you get through it by playing with it, being curious, and trying things out in a safe, ethical framework. Thank you so much — we're going to make sure we've landed on our learning objectives with that last slide, and then let's do questions.

Q&A Panel

KM: Yes. Let's see — I'm going to do my best to consolidate and compile the questions to fit the time we have.

Is it ethical to summarize case notes using AI?

KM: I wouldn't. And whoever left us this feedback, thank you — we can talk about it for a bit. Christy, would you like to go first?

KN: Yeah, I think for case notes, it's taking directly what you said in that voice message and helping you speed up a process using a template — this is the way I need to put the case note into a document, use this template. It takes the voice recording, puts it in, and then you take out things that don't make sense. It's not necessarily summarizing and pulling out key details — it's automating or speeding up a process that could otherwise be time-consuming. That's the hope we were trying to convey with that demonstration.

Which free AI tool is best for translations, or for something like flyers?

KM: Yep. I'm seeing a lot of questions about which platforms to use for various use cases. Again, if you're part of a larger organization, ask what you're approved to use and what's available — more often than not, there's something that solves the problem. And when you ask what's 'best,' it's really relative, because there are so many products doing almost identical work, and the technology is changing rapidly, so it's hard for me to say what's best. It's whatever your organization allows you to use, and whatever produces a result that's acceptable for you.

KN: And we can't endorse any products. I can actually tell you what—

KM: Yes.

KN: —you use for translation, so it really goes back to that AI readiness toolkit I mentioned at the start of the call, which we'll be sending out with more resources to continue your learning — those questions will help you identify the best tool as well.

Should we be deleting our data for cybersecurity reasons, and what are context windows or AI project spaces?

KN: Sarah has a great question about whether we should be deleting our data for cybersecurity reasons after use — I'd mentioned context windows and creating AI project spaces, which I use in Claude, to remember my context so it's easier to collaborate when I come back to it. This is kind of a two-part question, so Kim, jump in with more on that.

KM: Yes.

KN: It's not necessarily that I'm using PII or personal data when I'm designing those spaces — it's more like, here's my role, here's my goals, here's a lot of my daily tasks, and I'm training the AI feature to understand why I'm using it and what my goals are, so it can collaborate with me better. But I understand a cybersecurity concern would be more like, 'I did a data analysis on this data and I don't want it retained' — I always de-identify data before putting it into a Claude-type tool anyway, but I'll let Kim take that distinction.

KM: Yeah.

KN: For Sara's question.

KM: I realize we're almost at time, so I'll speak a bit fast. I noticed Krysti mentioned a term called context window, so I'll talk about it. There's a certain amount of memory a tool has before it forgets things — when you're 'training' AI over time, you're managing that context window, or its memory, and that varies by model. In this sense, we're talking about using a tool your organization has approved, so hopefully your organization or IT department has already established with the vendor what they do and don't retain — so deleting things stops being your responsibility. But you still want to be mindful of what you input, and use your best judgment if it's something sensitive, even with an approved tool. In this context, we're talking about things like task information, role information, and project objectives.

How do I safely use client data with AI tools?

KN: There was also a question about how to safely use client data. I wouldn't put client data into an AI tool — you can take out identifying information, or use placeholders like 'Client A' and 'Client B' when working on a document. Kim, would you take Maya's question? I think that's a good one for our last question.

What should you do if a client doesn't want to give informed consent for AI to be used in their case?

KM: Maya asks: what should you do if a client, or multiple clients, don't want to give informed consent for the use of an AI tool in their case? Is that grounds for them to no longer receive services, since it's a reasonable work tool, or should there be a channel where AI isn't used for clients who don't consent? That's a great question. They're owed the services no matter what tool you want to use — if they don't give consent, don't use it. That's my response from both an ethics and an efficiency perspective — it's just a tool, so you don't need it to do your job. It helps, but if you don't have consent, you can do the work without it, and you should.

KN: We just produced a guide on how to talk with clients about AI and ask for consent, because they need to understand what's being used before they can give it — giving that dignity and digital-literacy moment on what AI is, how it's used, and what it can and can't do, similar to this presentation's look at limitations and capabilities. It helps clients make informed decisions about how their data will be managed — sometimes clients will say yes if they know their data isn't being used for something like housing searches, partners, or community resources, so it really depends on the task and use case. Please look for that guide — we'll send it out in the follow-up email; it's a great resource for that question. We also saw some questions about things like PowerPoint and other design features — you can upload examples for that too, but it really depends on what your organization has approved, so once you know that, you can start building best practices around the tools that are allowable in your organization.

Conclusion

Reviewing Learning Objectives

KN: And with that, we hope today you now know how to describe generative AI, its capabilities, and its limitations; that you can identify at least two resettlement work tasks that AI could support; and that you can apply that three-step framework to determine if AI is suitable for a resettlement task. I appreciate all your honest answers and feedback — we appreciate your time. Here's the list of resources we'll send out; it's all about curiosity, following through, and doing what feels right for you and ethically sound for your organization and the clients we serve. We also want to make sure you take our five-question survey at the end, just to help us know how we did and how we can improve trainings like this in the future. Switchboard is here to support you all in the work you do, with resources, trainings, webinars, guides, toolkits, blogs, and technical assistance. Please take out your phones one more time — take a picture of the QR code, give us just a few more minutes for that feedback. Once that's done, we'll send a link to our website so you can explore the existing AI resources that Kim, I, and our other colleagues have put together — and there's that link for our website and for other TA requests. Thank you so much, Vincent, for the behind-the-scenes production on this webinar — we didn't see you, but we needed you, and you put in so much time and effort, so a shout-out to Vincent.
KM: Yes, thank you.

The IRC received competitive funding through the U.S. Department of Health and Human Services, Administration for Children and Families, Grant #90RB0053. The project is 100% financed by federal funds. The contents of this document are solely the responsibility of the authors and do not necessarily represent the official views of the U.S. Department of Health and Human Services, Administration for Children and Families.