

Demystifying AI: Practical Applicability in Social Services

Objectives:

- Ground a shared definition of “AI”
- Share current State approach, tools, and guardrails
- Review best practices for identifying strong AI use-cases
- Discuss AI applications in the context of social services



Agenda/Topics

- Defining “Artificial Intelligence”
- Overview of AI at SoV
- Identifying practical use cases (and risks)
- Broader concerns: environmental & workforce impacts
- AI in social services:
 - What AI does well
 - AI-specific risks
 - Equity & vulnerable populations
 - High vs. low-stakes use-cases



What is “Artificial Intelligence?”

- **1930s:** Robotics/Cybernetics/Photocells
- **1950s:** Neural Networks (on paper)
- **1960s-70s:** Machine Translation and Signal Processing
- **1980s-90s:** Executive Decision Systems/Decision Trees
- **2000s:** Machine Learning/Genetic Algorithms
- **2010s:** Computer Vision/Neural Networks
- **2020s:** Robotic Process Automation
- **2023:** Generative AI/GPT


“Computers doing stuff we thought took a human”



What “AI” means today

VT: Any algorithm, including one incorporating machine learning or other artificial intelligence techniques, that uses data-based analytics to make or support government decisions, judgments, or conclusions (3 V.S.A. § 3305).


- **Practically:** statistical pattern-finding, not “thinking” or “reasoning.”
- **What people mean now:** foundation models, especially LLMs, used via chat, plug-ins, or embedded in apps.



The State's current approach

- Principles-first: Vermont's [AI Code of Ethics](#) and [AI Use Guidelines](#).
- Core posture:
 - **Human-Centered Use:** *Respect individual dignity, enhancing service delivery without infringing on freedoms.*
 - **Transparency and Accountability:** *Ensure clear documentation and human accountability for decisions.*
 - **Ethical Implementation:** *Support human-led processes, avoid autonomous decisions with significant impacts, continually monitor for fair outcomes.*
- **In short:** AI is a “power tool” for humans to responsibly improve the efficiency & quality of their work.

NOT IDEAL:




RECOMMENDED:



What's in use today across the State


- **ChatVT** for safe, general drafting/brainstorming and Q&A.
 - RAG knowledge repositories: agency docs + search/Q&A for internal use.
- **Document processing:** Document Intelligence for PDFs/forms.
- **Contact centers:** piloting Amazon Q/Connect assistants.
- **M365 Copilot:** pilot now; broader rollout targeted for mid-2026.
- **AI coding tools:** Code generation, rapid prototyping
- **Common wins:** Document parsing, first drafts, summaries, FAQs/knowledge repos.



Practical Risks



- **AI hallucinations**
 - Models not capable of knowing “truth”
- **Data loss/exposure**
 - Particular risk with public tools
 - Increases further with rise of embedded systems
 - Data governance is critical
 - **AI doesn't use the data it's supposed to access, it uses the data it can access.**



“According to the board that hears assessment appeals, Ren's arguments are riddled with fictitious case law — possible artificial intelligence (AI) ‘hallucinations’ that sent B.C.'s assessment authority on a wild goose chase”



Identifying AI Use-Cases

1. Start with ethical alignment and public value.
2. Define a specific, measurable problem.
3. Match tasks to AI strengths.
4. Understand and optimize human processes.
5. Assess risk, complexity, and data needs early.





Common Implementation Risks

- **Data sensitivity**
 - AI tools may expose confidential data during processing
- **Integration Complexity**
 - Solutions spanning multiple systems can fail due to API gaps, identity management issues, etc.
- **Reliability & Bias**
 - AI outputs may be inaccurate, unreliable, or biased
- **Operation & Governance**
 - Lack of documentation on process/ownership can lead to ambiguity



The shift: embedded and agentic AI

- From **chatbots** to **systems**:
 - **Embedded AI**: shows up inside apps (email, CRM, case mgmt.).
 - **Agentic AI**: plans multi-step tasks and uses tools/APIs automatically.
- **New risks**:
 - Hidden model use; quiet data movement via connectors
 - Automation surprises; permission creep; provenance gaps
- **What changes for us**:
 - Must evaluate systems, not just chat UIs.
 - Stronger change management and monitoring.
 - **Education**: prescriptive approach is challenging




AI and the Environment

- AI data center water and energy usage forecast to increase ~10x over decade
- Areas with AI infrastructure buildup often have pre-existing water constraints
- Main energy/power draw is during model training phase
- **What to do:**
 - Ensure AI use is focused on efficiency
 - Minimizing cost correlates well with reducing environmental impact
 - Trust your ethical/value judgement!



All data centers combined use as much power as some of the world's largest economies

Electricity demand 2023; thousands of terawatt-hours




Sources: International Energy Agency; Organization of the Petroleum Exporting Countries; and IMF staff calculations. Note: Electricity demand for data centers compares with that in biggest natural users as of 2023. EVs = Electric vehicles.

IMF

Data centers electricity consumption set to more than double by 2030


Electricity consumption will surge due to the high energy demands of artificial intelligence, in data centers mainly located in the US and China



Source: International Energy Agency

AFP


Global Data Center Water Consumption



Source: Bluefield Research. As of June 2023.



Data Center Infrastructure in the United States, 2025




AI and the Workforce

- AI impact us vastly different than previous “industrial revolutions”
- Most impacted job categories include mathematics, computer science, etc.
- Jobs will not simply “disappear”, but impact is hard to predict
- Computer: was a job in 1940s, then became a system
 - Dozens of different job types developed to support



Chart 1 AI scenarios

1990 dollars (thousands), log scale




NOTES: The blue line is real gross domestic product (GDP) per capita in 1990 dollars. The orange line is a trend line fitted to the data for 1870-2024 with a trend growth rate of 1.9 percent per year. The red, green and purple lines are hypothetical paths for per capita GDP based on different scenarios.

SOURCES: Bureau of Economic Analysis; Haver Analytics; Macrohistory.net; United Nations; authors' calculations.

Federal Reserve Bank of Dallas



Jobs and Average AI Potential by Occupational Group



McInerney, Mark (April 11, 2025). "AI Task Force: Implications for Maine's Workforce," *Maine Center for Workforce Research and Information*.

AI in Social Services: What We'll Cover

What does AI look like in your day-to-day work, and what do you need to know to use it safely?

- ▶ What AI does well in social work
- ▶ AI-specific risks for your clients and practice
- ▶ Equity and vulnerable populations
- ▶ High- vs. low-stakes use cases and practical guardrails
- ▶ Real-world examples
- ▶ The human irreplaceable: reframing AI and your role



What AI Does Well in Social Work

- **Administrative burden reduction**
 - Business process optimization: invoice processing, etc.
 - Case note drafting from worker verbal summaries or transcripts
 - Meeting and call transcription (reducing manual note-taking)
 - Summarizing lengthy records: court orders, medical history, prior case files
- **Communication and access**
 - Real-time translation for multilingual client interactions
 - Plain-language rewrites of complex benefit or legal notices
- **Information retrieval**
 - Searching internal policy/procedure repositories (RAG tools)
 - FAQ bots helping clients navigate benefit systems



AI-Specific Risks: Confidentiality

- **Client data entered into public AI tools may be stored, logged, or used for model training**
 - Never enter names, case IDs, DOBs, diagnoses, or other PII into ChatGPT, Copilot (personal), or any other public/unapproved tool
 - Use only secure tools for any client-adjacent work
- **HIPAA and confidentiality obligations apply regardless of the tool used**
 - Users are responsible for understanding AI tool data use and privacy policies
- **Embedded AI is harder to see — and harder to control**
 - AI is increasingly built into apps you already use (Teams, Outlook, forms systems)
 - Know which tools in your environment have AI features enabled



AI-Specific Risks: Hallucination in Case Contexts

- **AI models cannot verify truth, they generate plausible-sounding text.**
 - A case note drafted by AI may contain invented details, dates, or quotes that never occurred
 - Once entered into a case file, hallucinated content becomes “official record”
- Real-world example: Canadian legal case where AI-generated filings cited fabricated case law
 - The same risk applies to case histories, eligibility rationales, and safety assessments
- **Rule: Never paste AI output directly into a case file without careful human review**
 - Treat AI drafts like a first draft from a new employee. Always evaluate and verify.

[Reality Check: Alberta Court Weighs Costs Award Against Counsel for AI-Generated Case Law Hallucination | Blakes](#)



AI-Specific Risks: Algorithmic Bias

- **AI systems trained on historical data can encode and amplify existing inequities**
 - Predictive risk tools in child welfare have shown documented racial disparities
 - Facial recognition and behavioral scoring tools perform worse on people of color
- **This is not theoretical:** Allegheny County (PA) risk scoring tool flagged Black families at higher rates, sparking national scrutiny
- **What to watch for in your setting**
 - Any tool that produces a “risk score” or recommendation for a client action
 - Ask: Who built this model? What data trained it? Has it been audited for bias?
 - A score is a data point, not a decision — human judgment must remain primary

[Evaluation Findings on the Use of Predictive Risk Models in Child Welfare - Allegheny Analytics](#)



AI-Specific Risks: Over-Reliance and Client Consent

- **Over-reliance reduces the quality of professional judgment over time**
 - Workers who default to AI summaries may miss nuance visible only in the full record
 - Safety assessments, removal decisions, and service plans require full human engagement
 - **Oversight compression:** “human in the loop” in theory vs. practice
- **Clients may not know AI is involved in their case**
 - Do clients know their calls are being transcribed? That a tool is scoring their risk?
 - Vermont’s AI Ethics principles emphasize transparency: consider what disclosure looks like in your program
- **AI cannot replicate human relationships or judgement**
 - Trust, rapport, and trauma-informed presence are not automatable



Equity and Vulnerable Populations

- **The digital divide affects your clients directly**
 - AI-facing systems (chatbots, online portals) assume internet literacy, internet access, tech comfort
 - Clients with disabilities, low literacy, or limited English may be further excluded
- **Language model limitations with non-dominant English**
 - AI performs worse on AAVE, regional dialects, and non-native English
 - Culturally-specific language may be misread or flattened
 - Translation & AI captioning cannot be used in accuracy critical applications without human review
- **Systemic bias can be baked into automation**
 - Benefit eligibility screening tools and triage systems may disadvantage marginalized clients
 - Automation at scale means errors affect thousands, not just one person
- **Advocacy opportunity:** social workers are well-positioned to flag inequitable AI outcomes



High- vs. Low-Stakes Use Cases

- **LOW STAKES: AI as drafting assistant (appropriate with review)**
 - First draft of case notes from your verbal summary
 - Translating a letter for a client into another language
 - Summarizing a policy document for your own understanding
 - Brainstorming community resources or service options
- **HIGH STAKES: AI must not drive the decision (human primacy required)**
 - Safety assessments and removal or placement decisions
 - Eligibility determinations with significant benefit or legal consequences
 - Any output that goes directly into a legal or regulatory record without review
 - Risk scoring that informs supervision level, court recommendations, or service denial
- **When in doubt:** the higher the stakes for the client, the more human oversight is required



Practical Guardrails for Workers

- **Never enter client PII into unapproved AI tools; use only secure/approved platforms**
- **Review all AI-generated content before it enters a case file, email, or report**
 - Ask “Is this accurate? Can it be verified?”; not just “Does this sound right?”
 - If AI produces something you can’t verify, don’t depend on it
- **Disclose AI involvement when it is material to a client interaction or decision**
- **Escalate concerns about AI tools in your systems**
- **Document your process:** if AI assisted a task, note how it was used and reviewed
- **Remember:** You are accountable for your work product, regardless of how it was generated



Real-World Examples in Social Services

- **Case note generation: Pilots underway in multiple state child welfare agencies**
 - Workers dictate or narrate; AI drafts the note; supervisor and worker review before filing
 - Early results: reduced documentation time, but quality review is essential
- **Translation: Refugee resettlement and immigration services**
 - Real-time translation tools improving access for clients with limited English
 - Still requires bilingual staff review for accuracy in sensitive conversations
- **Benefits navigation: AI chatbots helping clients understand eligibility**
 - Reduces call volume for simple questions; frees workers for complex cases
- **Crisis triage: 988 and crisis line augmentation tools**
 - AI surfacing resources and de-escalation prompts for counselors in real time
 - Human counselor always remains in control and accountable



How to Evaluate AI Tools

- **Focus on:** Privacy policy, regulatory compliance
- **Questions to ask:**
 - Data handling & storage
 - Regulatory compliance
 - Access & accountability
 - Bias & fairness
- **The baseline rule:** If a vendor can't clearly answer questions about data storage, HIPAA status, or who accesses your data, that's your answer.



The Human is Irreplaceable

- **AI will not replace social workers, but it will change the job.**
 - The administrative burden AI can reduce is real; that reduction is genuinely valuable
 - Workers freed from paperwork can spend more time on relationships, advocacy, and direct care
- **What AI cannot do:**
 - Build trust with a client in crisis
 - Read a room, notice body language, respond to grief
 - Exercise ethical judgment in ambiguous, high-stakes situations
 - Advocate for a client whose situation falls outside the algorithm
- **Your professional judgment, values, and relationships are, and will remain, vital.**
- **AI is just a tool.**



Questions + Further Resources

- [Artificial Intelligence in Vermont | Agency of Digital Services](#)
- [The AI Social Worker](#)
- <https://www.revolutioninai.com/2026/03/free-ai-tools-for-social-workers-2026.html>

