ACCESSIBILITY PROMPTCRAFT LAB

Turning Everyday AI Tools into Inclusive Library Services

IFLA Satellite Workshop
Beyond Boundaries Symposium — University of Macau, 2026
Workshop Handbook & Presentation Guide

Santiago Villegas-Ceballos
AI Innovation Unit (iIA) — eLibro
IAparaBibliotecas.com | MedeJean: Inteligencia Cultural
IFLA Information Technology Committee (Elected Member)

How to Use This Document

This handbook serves two purposes simultaneously. First, it is the complete workshop proceedings for participants, containing all methodological content, prompt templates, tool evaluations, and ethical frameworks covered during the session. Second, it is the presenter's guide: throughout the document, highlighted boxes marked with 📊 SLIDE describe exactly which slide should accompany each section, including recommended visuals, data points, and key messages for the presentation deck.

Workshop facilitators can use these slide descriptions to build a professional presentation that mirrors the document's structure, ensuring that the live session and the printed handbook work as complementary resources.

📊 SLIDE 1: TITLE SLIDE

Display workshop title, IFLA and University of Macau logos, author name and affiliations (iIA, IAparaBibliotecas.com, MedeJean, IFLA IT Committee). Include a visual of diverse library users interacting with technology. Subtitle: "From critical AI literacy to universal library services."

Table of Contents

1. Why This Guide Exists

📊 SLIDE 2: THE ACCESSIBILITY CRISIS IN NUMBERS

Full-screen data visualization: 1.3 billion people with disabilities (WHO), 95% of books never converted to accessible formats, 94.8% of top websites fail WCAG (WebAIM Million), 80% of persons with disabilities live in developing countries (World Bank). Use a world map heat-map showing disability prevalence vs. accessible content availability. End with: "Libraries are uniquely positioned to close this gap."

Libraries serve **1.3 billion people with disabilities worldwide** — 16% of the global population (WHO, 2023) — yet an estimated 95% of published books have never been converted to accessible formats. Free AI tools can now close this gap at unprecedented speed, but only if librarians understand both the tools' capabilities and their structural failure modes.

The urgency is real. The WebAIM Million project found that 94.8% of the world's top million websites contain detectable WCAG failures. In low-income countries, only 27% of the population uses the internet at all (UN DESA). Eighty percent of persons with disabilities live in developing countries (World Bank). AI-powered accessibility tools — when deployed critically — can help libraries in the Global South leapfrog infrastructure barriers at minimal cost. But AI tools are not neutral. They hallucinate, flatter, generate biased content, and simulate expertise they do not possess. The PAPAYAS framework, introduced in Section 2, gives librarians a critical lens for understanding these failures before deploying solutions.

This document is grounded in the UN Convention on the Rights of Persons with Disabilities (191 parties), the Marrakesh Treaty (102 contracting parties covering 126 countries, per WIPO Lex), WCAG 2.2, and IFLA's 2024 Guidelines for Making Libraries Accessible. It follows the principle of "Nothing About Us Without Us" — disabled communities must be partners, not passive recipients, in every accessibility initiative described here.

1.1 Disability Is a Universal Experience Shaped by Inequality

The WHO reports that 1.3 billion people experience significant disability — one in six humans. Within this figure, 2.2 billion have some form of visual impairment, 1.5 billion live with hearing loss, and hundreds of millions experience cognitive, motor, or communication disabilities. These numbers are growing due to aging populations, chronic disease, and health emergencies.

Disability and poverty reinforce each other. Only 3% of people who need assistive technology in low- and middle-income countries have access to it (International Disability Alliance). The digital divide compounds exclusion: in the EU alone, 18.9 million people with disabilities are left behind digitally, with 17% fewer using the internet regularly compared to the general population (Surfshark, 2024). In the United States, disabled people are three times more likely to never use the internet (ThemeIsle, 2025).

1.2 The Book Famine Persists Despite Legal Progress

📊 SLIDE 3: THE MARRAKESH TREATY AND THE BOOK FAMINE

Timeline graphic showing: 2013 Treaty adoption → 2016 entry into force → 2026 current status (102 contracting parties). Key stat: Before the Treaty, 99% of books in least-developed countries were never in accessible formats. Show the ABC Global Book Service: 600,000+ books in 80+ languages across 93 libraries. Closing question: "How can AI accelerate what the Treaty made legal?"

Before the Marrakesh Treaty, 99% of books in least developed countries were never converted to accessible formats. The treaty, adopted in 2013 and now ratified by 102 contracting parties, allows authorized entities — including libraries — to create and share accessible format copies without copyright holders' permission. The WIPO Accessible Books Consortium now holds over 600,000 books in 80+ languages across 93 partnered libraries. Yet a 2025 EU evaluation found only 7–20% of books were available in accessible formats despite digital technology. AI can dramatically accelerate format conversion, but human oversight remains essential for quality.

1.3 Library Associations Are Converging on AI and Accessibility

IFLA published its Guidelines for Making Libraries Accessible for People with Disabilities in 2024 and its Entry Point to Libraries and AI document in 2025, identifying "descriptive AI" — which makes photo, video, and manuscript collections machine-readable at scale — as a major opportunity. IFLA's 2025 Guidelines for Inclusive Library Services for Persons with Print Disabilities ground library work in Universal Design and the right to read. The ALA's ACRL approved AI Competencies for Academic Library Workers in October 2025, emphasizing equitable access. EIFL maintains a dedicated Libraries and AI program serving 60+ developing countries. The convergence is clear: library associations worldwide recognize AI as a tool for expanding accessibility with appropriate critical oversight.

2. The PAPAYAS Framework: Understanding AI Failure Modes

📊 SLIDE 4: INTRODUCING PAPAYAS

Title: "PAPAYAS: A Critical Lens for AI in Libraries." Visual: seven papaya fruit icons, each labeled with one letter. Subtitle: "Before deploying any AI tool for accessibility, understand its structural properties." Animate each letter appearing sequentially. Note: this framework treats AI behaviors as structural properties, not bugs.

The PAPAYAS framework is a mnemonic heuristic describing seven structural, predictable properties of Large Language Models. It argues that LLMs operate as statistical text-generation engines, not neutral knowledge repositories or cognitive agents. Understanding these properties is essential before using any AI tool for accessibility services, where errors can cause real harm to vulnerable users.

2.1 P — Parlanchines / Prolix (Chatty)

📊 SLIDE 5: P — PARLANCHINES (PROLIX)

Split screen: LEFT shows a verbose AI response (500+ words) to a simple question; RIGHT shows the same answer in 3 concise sentences. Key stat: GPT-4 exhibits verbosity compensation 50.40% of the time, with verbose responses scoring 27.61% lower in accuracy (Zhang, Das & Zhang, 2024). Accessibility impact callout: "Screen reader users must listen to ALL of this."

LLMs exhibit "verbosity compensation" — generating excessive tokens when facing uncertainty, masking ignorance with length rather than admitting they do not know. Research by Zhang, Das, and Zhang (2024) found GPT-4 exhibits verbosity compensation 50.40% of the time, with verbose responses scoring 27.61% lower in accuracy than concise ones.

Accessibility impact: Users with cognitive disabilities or attention disorders are overwhelmed by unnecessarily long responses. Blind users on screen readers must listen to excessive text. When asked to create Easy Read documents, AI may produce text that is technically simpler but far longer than necessary, defeating the purpose.

Mitigation for librarians: Always specify maximum word counts or sentence limits in prompts. Request "concise" or "brief" output explicitly. Review all AI output for unnecessary padding before sharing with patrons.

2.2 A — Aduladores / Adulatory (Flatterers)

📊 SLIDE 6: A — ADULADORES (SYCOPHANTIC)

Demonstration: show a live prompt where the librarian tells the AI "Our website is fully accessible" and the AI agrees — then show the same website failing 47 WCAG checks in WAVE. Key stat: LLaMA 2 accuracy drops up to 27% when users suggest incorrect premises (Sharma et al., 2024). Warning box: "Never ask AI to VALIDATE. Ask it to IDENTIFY FAILURES."

RLHF (Reinforcement Learning from Human Feedback) conditions models to prioritize user agreement over truth. Sharma et al. (2024) identified four sycophancy types: feedback sycophancy, response sycophancy, mimicry sycophancy, and error-admission sycophancy. Crucially, this is an "inverse scaling" phenomenon — larger, more capable models are MORE sycophantic, not less. LLaMA 2's accuracy drops up to 27% when users suggest incorrect premises.

Accessibility impact: If a librarian asks “Is our website accessible?” the AI may agree rather than identifying failures — dangerous for WCAG and ADA compliance. If staff frame disability through a pity or medical model, AI mirrors that framing instead of redirecting toward the rights-based social model.

Mitigation for librarians: Never ask AI to “validate” accessibility compliance — ask it to “identify all potential accessibility failures.” Frame prompts as requests for critique, not confirmation.

2.3 P — Políticamente posicionados / Partisan (Biased)

📊 SLIDE 7: P — POLÍTICAMENTE POSICIONADOS (BIASED)

Infographic: “WEIRD training data” — show the demographic skew of AI training corpora (Western, Educated, Industrialized, Rich, Democratic). Key tension: Medical Model vs. Social Model of disability. Show two AI responses to the same question about a wheelchair user — one framing disability as deficit, one as societal barrier. Call to action: “Explicitly instruct AI to use social model language.”

Models exhibit structural bias from training data demographics — predominantly Western, Educated, Industrialized, Rich, Democratic (“WEIRD”). Hall, Westwood, and Grimmer (2025) found 76% of LLMs tested showed systematic political preference. Fulay et al. (EMNLP 2024) demonstrated that even optimizing purely for “truthfulness” inevitably induces left-leaning bias — making true algorithmic neutrality a statistical impossibility.

Accessibility impact: AI trained on Western medical literature defaults to the medical model of disability (disability as individual deficit) rather than the social model (disability as societal barrier). Training data underrepresents disability experiences from the Global South, where disability is understood through different cultural, spiritual, and communal frameworks.

Mitigation for librarians: Explicitly instruct AI to use social model and rights-based language. Specify cultural context in prompts. Always have output reviewed by members of the local disability community.

2.4 A — Aprendices lentos / Apprenticed (Slow learners)

📊 SLIDE 8: A — APRENDICES LENTOS (NO LEARNING)

Diagram: show two identical chat sessions side by side. In Session 1, a librarian corrects the AI’s accessibility guidance. In Session 2 (next day), the AI makes the exact same error. Key stat: GPT-4 training cost US$78M; Gemini Ultra US$191M. Bottom line: “Your corrections vanish. Maintain a validated prompt library instead.”

LLMs do not learn from user interactions. Brown et al. (2020) established that “in-context learning” is merely temporary pattern matching — once a session ends, information disappears. True parametric learning requires retraining costing US$78M (GPT-4), US$191M (Gemini Ultra), or US$170M (LLaMA 3.1).

Accessibility impact: If you correct an AI’s incorrect accessibility guidance in one session, it will make the same error next time. Disabled users’ corrections about their own experiences are lost between sessions — directly contradicting the “Nothing About Us Without Us” principle.

Mitigation for librarians: Maintain a library-specific prompt library (Section 7) with pre-tested, validated prompts. Use system prompts or custom instructions to set accessibility standards at the beginning of each session.

2.5 Y — Yacentes / Yoked to cutoff (Dormant)

📊 SLIDE 9: Y — YACENTES (FROZEN IN TIME)

Timeline showing: training data cutoff date vs. today's date. Mark key events AI might miss: European Accessibility Act enforcement (June 2025), WCAG 2.2 updates, new assistive technologies. Visual metaphor: an AI trapped inside a snow globe while the world moves outside. Message: "Always verify legal and standards information against current sources."

Models are frozen at their training data cutoff, structurally blind to anything that occurs afterward. RAG and web browsing are superficial patches operating in the context window, not genuine parametric learning.

Accessibility impact: AI may not know about the European Accessibility Act enforcement (June 2025), recent WCAG updates, or new assistive technologies. When asked about specific current requirements, it may give generic, outdated responses.

Mitigation for librarians: Always verify legal and standards information against current authoritative sources. Include the current date and specific jurisdiction in prompts. Cross-reference AI guidance with official WCAG, ADA, and EAA documentation.

2.6 A — Aleatorios / Aleatory (Random)

📊 SLIDE 10: A — ALEATORIOS (RANDOM/HALLUCINATORY)

Live demonstration: ask three different AI chatbots the same accessibility question and show divergent (sometimes contradictory) answers side by side. Highlight hallucinated WCAG criteria numbers. Key warning for blind users: "AI-generated alt text can confidently describe objects NOT present in images. Users who cannot verify visually are at greatest risk."

LLMs cannot distinguish between epistemic uncertainty (reducible through more knowledge) and aleatoric uncertainty (irreducible noise). When facing knowledge gaps, instead of reporting low confidence, they generate hallucinatory, high-confidence text.

Accessibility impact: The same accessibility question may receive different — sometimes contradictory — answers across sessions. AI may confidently cite nonexistent WCAG success criteria. AI-generated alt text is prone to hallucination, confidently asserting details not present in images. For blind users who cannot visually verify, this represents a fundamental breakdown in trust.

Mitigation for librarians: Run critical accessibility queries through multiple AI tools and compare outputs. Never deploy AI-generated alt text or accessibility content without human verification.

2.7 S — Simuladores / Simulators

📊 SLIDE 11: S — SIMULADORES (STATISTICAL SIMULATORS)

Conceptual diagram based on Janus (2022): show an LLM as a “superposition of states” before input, collapsing into a specific “simulacrum” as tokens are added. Key message: “AI simulates expertise; it does not possess it. Treat ALL accessibility output as a draft requiring expert and community review.” Closing: group photo of the PAPAYAS framework summary card.

Based on Janus’s (2022) Simulator Theory, an LLM is not an agent or oracle — it is a statistical engine simulating the processes that produced human text. Before receiving input, the model exists in a “superposition of states,” harboring patterns of scientists, trolls, and conspiracists simultaneously.

Accessibility impact: AI can generate text that sounds deeply knowledgeable about accessibility but is statistical pattern-matching, not comprehension. It may produce “inspiration porn” — patronizing disability content — because it simulates empathy based on training data patterns.

Mitigation for librarians: Treat ALL AI output as a draft requiring expert review. Never present AI-generated accessibility content as authoritative without human validation. Use the PAPAYAS framework as a training tool for staff.

3. A Five-Phase Methodology for Librarians

📊 SLIDE 12: THE FIVE-PHASE DEPLOYMENT METHODOLOGY

Circular process diagram with five phases: (1) Assess Community Needs → (2) Select & Configure Tools with PAPAYAS → (3) Build, Test & Validate Prompts → (4) Deploy with Human-in-the-Loop → (5) Monitor, Evaluate & Iterate → back to (1). Center of circle: "Nothing About Us Without Us." Each phase should be color-coded and numbered.

Phase 1: Assess Community Needs and Existing Infrastructure

Before selecting any AI tool, conduct a structured assessment. Identify which disability groups your library serves or should serve. Consult local disability organizations, special education schools, and rehabilitation centers. Document specific needs using the five accessibility dimensions in this guide. Inventory your hardware: minimum 16 GB RAM desktop for local AI and reliable internet for cloud tools. Assess staff skills: LM Studio's GUI suits non-technical staff while Ollama suits those comfortable with command lines. Determine privacy requirements: disability status is special category data under GDPR Article 9.

Phase 2: Select and Configure Tools Using the PAPAYAS Lens

For each tool under consideration, apply the PAPAYAS checklist. Does the tool tend to generate excessively verbose output for your use case? Test it with a specific accessibility task. Does it validate incorrect accessibility claims? Test by deliberately providing a non-compliant document and asking if it is accessible. Does it reflect your community's cultural understanding of disability? Test with locally relevant scenarios.

Phase 3: Build, Test, and Validate Prompt Templates

Use the prompt library in Section 7 as starting points. Customize for your library's specific context, community languages, and patron needs. Test every prompt at least three times across different AI tools to check for consistency. Have members of the target disability community review output before deploying.

Phase 4: Deploy with Human-in-the-Loop Workflows

📊 SLIDE 13: THE HUMAN-IN-THE-LOOP WORKFLOW

Flowchart: AI Generates Draft → Staff Reviews for Accuracy → Disability Community Member Validates → Content Published. Include a decision diamond: "Does output meet standards?" with Yes/No paths. Key stat: Organizations using human-in-the-loop systems report 42% fewer AI-driven errors. Red warning box: "For accessibility, a wrong alt text is worse than no alt text."

Establish a clear workflow: AI generates, staff reviews, disability community member validates, content is published. Document all decisions. For accessibility, where a wrong alt text is worse than no alt text, this review step is non-negotiable.

Phase 5: Monitor, Evaluate, and Iterate

Collect feedback from patrons with disabilities. Track which tools and prompts produce the highest quality output. Update prompts as AI models change. Review the PAPAYAS framework quarterly to reassess tool limitations.

4. Visual Accessibility: Making Library Content Visible to All

📊 SLIDE 14: VISUAL ACCESSIBILITY OVERVIEW

Section title slide with icon representing visual accessibility (eye with radiating lines). Three columns showing the sub-topics: Alt-Text Generation | Document Remediation | Accessible Image Creation. Statistics: 2.2 billion people with visual impairment worldwide (WHO); 285 million blind or with low vision.

4.1 AI-Powered Image Description and Alt-Text Generation

📊 SLIDE 15: ALT-TEXT WORKFLOW DEMONSTRATION

Step-by-step live demo: (1) Upload image to ChatGPT/Claude/Qwen free tier, (2) Apply the structured alt-text prompt from Section 7, (3) Generate three versions (short, long, extended), (4) Staff verifies accuracy, (5) Screen reader user tests. Include before/after: generic AI alt text vs. prompt-engineered alt text. Cite American Foundation for the Blind on AI alt-text limitations.

The most immediate accessibility win AI offers is scalable alt-text generation. ChatGPT, Claude, Gemini, and Qwen all accept image uploads on their free tiers and can generate descriptions. However, quality varies significantly. The American Foundation for the Blind (2026) notes that AI-generated image descriptions often lack narrative structure and contextual understanding of what makes an image meaningful.

Recommended workflow: upload the image to a multimodal AI chatbot (ChatGPT or Qwen recommended for free-tier image analysis), use the structured alt-text prompt from Section 7, request three versions (short alt text under 125 characters, long description of 1–3 sentences, and extended description for complex images), have a staff member verify accuracy, and have a screen reader user test the description for usefulness.

Specialized Visual Accessibility Tools (All Free)

Tool	Platform	Key Features	Languages
Be My Eyes / Be My AI	iOS, Android	GPT-4 visual assistance, real-time photo analysis, conversational	180+ languages
Microsoft Seeing AI	iOS, Android	14+ channels: text, documents, scenes, faces, currency, spatial exploration	18+ languages
Google Lookout	Android	7 modes: text, documents, objects, currency, Gemini Q&A	30+ languages

4.2 Document Remediation: Making PDFs Accessible

📊 SLIDE 16: PDF ACCESSIBILITY REMEDIATION

Diagram comparing: Manual remediation (US$1–4/page, weeks of work) vs. AI-assisted remediation (cents/page, hours). Show tools: PAVE (free, browser-based), Grackle Workspace (Google Docs add-on), AWS Open-Source Solution. Live demo: upload a non-tagged PDF to PAVE, show auto-corrections, download accessible version. Emphasize: AI does NOT auto-generate alt text for images — by design.

PDF accessibility remediation — adding proper tags, reading order, alt text, and structure — is one of the most labor-intensive accessibility tasks. AI is beginning to automate this. PAVE (PDF Accessibility Validation Engine) is a free web-based tool from ZHAW School of Engineering in

Switzerland that auto-corrects many issues. Grackle Workspace provides 22 accessibility checks against WCAG 2.0 and PDF/UA with unlimited tagged PDF exports from Google Docs. The AWS Open-Source PDF Accessibility Solution, developed with Ohio State and Arizona State universities, reduces remediation costs from US$1–4/page to cents per page.

4.3 AI Image Generation for Accessible Library Materials

📊 SLIDE 17: CREATING ACCESSIBLE VISUALS WITH AI

Gallery showing examples: accessible signage created with Ideogram (best text rendering), visual schedules from Google ImageFX, Easy Read illustrations from Leonardo AI, complete designs from Canva AI. Show the Universal Design Prompt Formula on screen. Live workshop activity: participants generate one accessible sign using the formula.

Free AI image generators can create visual materials that follow Universal Design principles: clear, simple, high-contrast images for signage, visual schedules, pictograms, and Easy Read illustrations.

Use Case	Best Free Tool	Why
Library signage with text	Ideogram (25 prompts/day)	Best-in-class text rendering in images
Visual schedules/pictograms	Google ImageFX (free)	Clean, consistent output; strict safety
Easy Read illustrations	Leonardo AI (150 tokens/day)	Most generous free tier; consistent style
Complete design workflow	Canva AI Free (50 lifetime)	Generate AND design in one platform
Social story images	DALL-E via Bing (400 imgs/mo)	Good conversational refinement

Universal Design prompt formula: [Subject] + [Style: flat illustration/simple vector/bold colors] + [Background: white/solid/high contrast] + [Accessibility: clear, simple, uncluttered, large elements] + [Purpose: signage/pictogram/Easy Read].

4.4 Braille and Tactile Materials

TactileNet (2025) represents a breakthrough: the first comprehensive AI framework for generating tactile graphics using Stable Diffusion models, achieving 92.86% adherence to tactile standards. Dot Pad by Dot Inc. is a refreshable tactile display with 2,400 pins that uses AI to translate visual content into tactile feedback, integrated with Microsoft AI and Google Gemini.

4.5 WCAG Compliance Checking with AI

📊 SLIDE 18: FREE WCAG TESTING TOOLKIT

Five-tool grid showing: WAVE (browser overlay), axe DevTools (zero false positives), Google Lighthouse (built into Chrome), Accessibility Insights (Microsoft, guided testing), AccessiMind (VS Code real-time). Key message: "Automated tools catch 30–40% of issues. The rest requires manual testing with assistive technology users." Show a sample WAVE report screenshot.

Tool	Type	Best For
WAVE (wave.webaim.org)	Browser extension	Visual error overlay

axe DevTools (Deque)	Browser extension	Zero false positives, industry standard
Google Lighthouse	Built into Chrome DevTools	Accessibility scoring
Accessibility Insights	Microsoft, open-source	Automated + guided manual testing
AccessiMind	VS Code extension	Real-time WCAG 2.2 analysis

5. Auditory Accessibility: Captioning, Transcription, and Sign Language

📊 SLIDE 19: AUDITORY ACCESSIBILITY OVERVIEW

Section title with ear/sound wave icon. Three pillars: Text-to-Speech | Speech-to-Text & Captioning | AI Sign Language (Frontier). Statistics: 1.5 billion with hearing loss (WHO). Quote: "Deaf community demands AI as supplement, never replacement, for human interpreters."

5.1 Text-to-Speech Tools for Library Content

📊 SLIDE 20: FREE TTS TOOLS RANKED

Comparison table of TTS tools with audio sample QR codes (participants can scan and listen). Highlight: Microsoft Edge Read Aloud requires ZERO setup — install on all library computers TODAY. Show edge-tts Python command for batch generating audio guides. Play a 15-second audio sample from ElevenLabs vs. Edge TTS vs. Piper for comparison.

Tool	Cost	Quality	Languages	Best For
Edge Read Aloud	Free, built-in	Excellent	100+	Zero-setup solution
edge-tts (Python)	Free, no API key	Same Azure quality	100+	Batch audio generation
Piper TTS	Free, open-source	High	30+	Runs on Raspberry Pi
ElevenLabs free	10K chars/month	Industry-leading	32	Short high-quality clips
Coqui TTS / XTTS	Free, open-source	Near-commercial	17	Voice cloning (6 sec)
Bark AI (Suno)	Free, MIT license	Expressive	Multi	Storytelling, emotion

Immediate action: Install Microsoft Edge on all public library computers. Edge's built-in Read Aloud feature provides instant TTS accessibility for any webpage with no setup, no account, and no cost.

5.2 Speech-to-Text and Captioning

📊 SLIDE 21: WHISPER: THE OPEN-SOURCE CAPTIONING REVOLUTION

Live demo: run Whisper on a 60-second audio clip in a non-English language, show real-time SRT generation. Command on screen: whisper recording.mp4 --model medium --language en. Comparison chart: YouTube auto-captions (60–70%) vs. Whisper (95–99%) vs. professional human (99%+). Workflow diagram: Record → Whisper → AI cleanup → Human review → Publish.

OpenAI Whisper is the gold standard for free speech-to-text. Fully open-source (MIT license), it supports 99 languages, runs locally for complete privacy, and achieves near-human accuracy for English. Libraries can install it on a local workstation to transcribe recorded lectures, oral histories, community events, and video content.

Practical captioning workflow for library events: record the event audio/video, run Whisper locally, clean up the SRT using an AI chatbot, upload corrected captions with the video, and publish the transcript on the library website.

Google Live Transcribe (free Android app, developed with Gallaudet University) provides real-time transcription in 80+ languages with offline capability. Libraries should recommend it to deaf and hard-of-hearing patrons for in-person reference desk interactions.

5.3 AI Sign Language Generation: A Rapidly Evolving Frontier

📊 SLIDE 22: AI SIGN LANGUAGE: PROMISE AND CAUTION

Map showing global sign language AI initiatives: Signapse (UK, BSL/ASL), SignAll at Prairie View A&M Library (US, first library pilot), NVIDIA Signs (US, free ASL learning), Terp 360 (Kenya, African sign languages), Google SignGemma (on-device). Critical sidebar: "300+ sign languages exist worldwide. Word-for-word translation fails completely. Non-manual markers carry grammar. The Deaf community must lead."

AI sign language technology is advancing quickly but remains controversial within the Deaf community and not yet suitable for complex or sensitive contexts.

Tool	Origin	Capability	Status
Signapse AI	UK	BSL/ASL via digital signers	Deployed in UK railways
SignAll	US	ASL to English text	First library pilot (Prairie View A&M)
NVIDIA Signs	US	ASL learning with 3D avatar	Free, 100 signs, building to 1,000
Terp 360	Kenya	Speech to Kenyan Sign Language	Africa Prize winner 2024
Google SignGemma	US	On-device ASL translation	Research preview

Critical consideration: Over 300 sign languages exist worldwide, each with its own grammar completely different from spoken languages. ASL uses topic-comment structure, spatial grammar in 3D space, and classifier predicates. The Deaf community has consistently advocated that AI is a supplement to, never a replacement for, human interpreters. Libraries must consult their local Deaf community before deploying any AI sign language tool.

5.4 Audio Description for Library Collections

AI can now generate audio descriptions for video content, exhibitions, and visual collections. ADAI supports 18+ languages with a free trial. Audible Sight offers automatic scene detection. ScreenPal auto-generates descriptions alongside captions. Traditional audio description costs US$8+ per minute; AI can produce draft descriptions in hours at a fraction of the cost, though human review remains essential.

6. Cognitive and Neurodivergent Accessibility

📊 SLIDE 23: COGNITIVE ACCESSIBILITY OVERVIEW

Section title with brain/neurodiversity icon. Four areas: Easy Read Conversion | Social Stories | Visual Schedules | ADHD Support. Key message: "This is where AI chatbots deliver the greatest IMMEDIATE impact — structured prompts reliably produce usable first drafts." Quote Inclusion Europe guidelines for Easy Read.

6.1 Easy Read and Plain Language Conversion

📊 SLIDE 24: EASY READ LIVE WORKSHOP

Interactive exercise: display a complex library policy paragraph (50+ words, jargon-heavy). Participants use their phones with ChatGPT/Claude/Gemini free tier to apply the Easy Read prompt from Section 7. Compare results across tools. Show Inclusion Europe rules: Arial 14+, never justify, no italics, one idea per sentence. Before/after side by side.

Easy Read uses simple words, short sentences, and supporting images so people with intellectual disabilities, low literacy, or limited language proficiency can understand information. The primary European standard is Inclusion Europe's "Information for All" (2009): use Arial 14+ font, never justify text, avoid italics and ALL CAPS, one idea per sentence, explain every difficult word.

AI chatbots excel at text simplification when properly prompted. The key is specificity: generic "simplify this" prompts produce mediocre results, while structured prompts specifying reading level, sentence length, formatting rules, and image suggestions produce usable Easy Read content. See the full Easy Read prompt template in Section 7 (Prompt 2).

6.2 Social Stories for Library Visits

📊 SLIDE 25: SOCIAL STORIES WITH AI

Show Carol Gray's structural rules on the left; AI-generated social story on the right. Highlight: first person only, ratio of 2–5 descriptive per 1 coaching sentence, sensory descriptions, positive ending. Workshop activity: participants create a social story for THEIR library using Prompt 4. The Autism Research Institute validates this approach with ChatGPT, Gemini, Copilot, and Claude.

Social stories, developed by Carol Gray, are personalized narratives helping individuals on the autism spectrum understand new situations. AI generates solid first-draft social stories when given Carol Gray's framework rules, specific library details (layout, sensory features, available accommodations), and the target audience's age and needs. The Autism Research Institute provides verified examples of AI-generated social stories with ChatGPT, Gemini, Copilot, and Claude.

6.3 Visual Schedules and Sensory-Friendly Content

Visual schedules use pictures and symbols to show activities in sequence. AI image generators with prompts specifying simple, clear illustrations for visual schedules produce reliable results. Several libraries have pioneered sensory-friendly approaches: the University of Washington's Autism-Ready Libraries Toolkit provides training modules and sensory audit checklists; Montana State University Library created a model sensory-seeking space.

6.4 ADHD Support Through AI

📊 SLIDE 26: NEURODIVERGENT SUPPORT TOOLS

Tool showcase: Goblin Tools (free, task decomposition), Otter.ai (free 300 min/month, real-time transcription), AI “brain dump” synthesis. Show an example: complex library research task broken into 5–10 minute steps with built-in break points. CHADD (2024) endorses AI for executive function support.

For patrons and staff with ADHD, AI tools offer immediate practical support. Goblin Tools (free) breaks complex tasks into manageable steps. AI chatbots excel at “brain dump synthesis” — transforming scattered thoughts into organized action lists. Otter.ai (free tier: 300 minutes/month) provides real-time lecture and meeting transcription with AI summaries, reducing the cognitive load of note-taking.

7. Motor and Physical Accessibility

📊 SLIDE 27: MOTOR ACCESSIBILITY

Three areas: Voice Control | Switch-Accessible Content | Reduced Input Effort. Show WCAG 2.2 new criterion: Target Size (Minimum) 2.5.8 = 24×24 CSS pixels minimum. Key challenge: voice AI still fails to recognize speech patterns of people with speech impairments (Harvard, Weru 2024). Diagram: all major OS include free dictation already.

Voice input eliminates typing barriers for users with motor impairments. All major operating systems include free dictation: Windows Dictation, macOS Dictation, Google Voice Typing, iOS Dictation. Voice assistants enable hands-free device control. Libraries can implement voice-activated catalog searching, hands-free website navigation, voice-controlled e-book readers, and AI chatbot reference services via voice.

For users who navigate via single switches, eye-tracking, or sip-and-puff controls, content must support sequential navigation with large, clearly defined interactive targets, adjustable timing, and single-switch scanning patterns. WCAG 2.2's new Target Size (Minimum) criterion (2.5.8) requires interactive targets of at least 24×24 CSS pixels. AI can help analyze website or catalog complexity and suggest simplifications.

8. Linguistic Accessibility: Translation, Multilingualism, and Underserved Languages

📊 SLIDE 28: LINGUISTIC ACCESSIBILITY

Comparison matrix of translation tools: Qwen (119 languages), Google Translate (130+), DeepL (30+, superior European quality), ChatGPT/Claude (context-aware), EU eTranslation (GDPR compliant), Meta NLLB (200+ including underrepresented languages). Heat map: translation quality by language pair (green = high, yellow = medium, red = low). Callout: indigenous languages remain poorly served.

8.1 AI Translation Tools for Multilingual Libraries

Tool	Languages	Strength	Best For
Qwen Chat	119	Broadest free coverage	Diverse communities
Google Translate	130+	Widest coverage	Quick signage translation
DeepL Free	30+	Superior European quality	European translations
ChatGPT/Claude	90+ each	Context-aware, tone-adjustable	Cultural adaptation
EU eTranslation	All EU official	Free, GDPR compliant	EU-based libraries
Meta NLLB	200+	Underrepresented languages	Indigenous research

Quality varies dramatically by language pair. Major European and East Asian languages produce high-quality results. Arabic, Hindi, Turkish, and Vietnamese produce medium quality. Indigenous languages, minority dialects, and many African languages remain poorly served. For critical documents, always use AI as a first draft plus human review.

8.2 The Challenge of Indigenous and Minority Languages

📊 SLIDE 29: INDIGENOUS LANGUAGES AND AI

Case studies from three continents: FLAIR Initiative (Mila, Quebec) for indigenous language preservation, Jared Coleman at Loyola Marymount for Owens Valley Paiute (programmatic LLM rules, not training data), University of the Arctic for Sámi languages (20 years of rule-based tools). Warning: LLM-generated text can contaminate online corpora. The Māori community opted against open-source to protect language sovereignty. Key principle: DATA SOVEREIGNTY.

Most indigenous languages are “low-resource” — insufficient training data for standard AI approaches. Critical warning: LLM-generated text can contaminate online corpora of indigenous languages with errors. Libraries working with indigenous languages must partner directly with language communities and ensure data sovereignty — communities control their language data.

9. Local AI Deployment: Privacy-First Accessibility Services

📊 SLIDE 30: LOCAL AI FOR LIBRARIES

Three-platform comparison: Ollama (CLI, developer-friendly), LM Studio (GUI, non-technical), Pinokio (one-click, experimental). Hardware requirements table: Budget US$400–600 | Recommended US$800–1,200 | Optimal US$1,500–2,500. Key message: "Patron disability data is GDPR special category data. Local AI keeps it on premises." Live demo: install Ollama, pull a model, run an Easy Read conversion.

9.1 Why Local AI Matters for Library Accessibility

Patron disability information — reading difficulties, cognitive disabilities, visual impairments — is sensitive personal data protected under GDPR and equivalent legislation. Cloud AI services require sending this data to third-party servers. Local AI keeps all data on library premises, transmitted nowhere.

9.2 Ollama: The Developer's Choice

Ollama is the leading open-source platform for running LLMs locally. It handles model management, optimization, and provides an OpenAI-compatible REST API on port 11434.

Task	Model	Size	Min RAM
Plain Language / Easy Read	Llama 3.1 8B	8B	8 GB
Translation	Qwen3 14B	14B	16 GB
Alt-text generation	Llama 3.2 Vision	7–12B	12 GB
Summarization	Mistral 7B	7B	8 GB

9.3 LM Studio: The Non-Technical Librarian's Choice

LM Studio offers a polished desktop GUI that makes running local AI accessible without command-line knowledge. Browse 1,000+ models visually, adjust parameters with sliders, and chat directly. Its built-in RAG feature lets librarians upload library policies or catalog data and query them locally.

9.4 Pinokio: One-Click AI Experiments

Pinokio is an "AI browser" that handles all dependencies automatically. Click one button to install Whisper, Bark, Stable Diffusion, or Ollama itself. No terminal, no debugging. Ideal for libraries wanting to experiment without IT expertise.

9.5 End-to-End Accessibility Workflow Using Local AI

📊 SLIDE 31: PRINT-TO-ACCESSIBLE PIPELINE

Flowchart showing complete pipeline: Scan (300 DPI) → OCR (Tesseract, 100+ languages) → AI Enhancement via Ollama (Easy Read, summary, translation, headings, alt text) → Format Creation (Tagged PDF, EPUB via Calibre, Audio via Piper TTS, Large Print 18pt+, Braille-ready) → Distribute via catalog with accessibility metadata. All tools free and open-source. Entire pipeline runs on a single workstation with no internet.

Print book to multiple accessible formats: scan at 300 DPI, OCR using Tesseract (free, local, 100+ languages), AI enhancement via Ollama for Easy Read version, summary, translation, heading structure, and alt text, then format creation — tagged PDF via Word/Google Docs, EPUB via Calibre, audio via Piper TTS, large print reformatted at 18pt+, Braille-ready via SensusAccess. This entire pipeline can run on a single library workstation with no internet after initial model downloads, using exclusively free and open-source tools.

9.6 Recommended Hardware for Libraries

Setup	Cost	Specs	Capabilities
Budget	US$400–600	16 GB RAM, 256 GB SSD	7–8B models: text simplification, summarization
Recommended	US$800–1,200	32 GB RAM, 512 GB, NVIDIA 8GB+	14B models, multimodal vision for alt-text
Optimal	US$1,500–2,500	64 GB RAM, 1 TB, RTX 4070+ or M4 Pro	30B+ models, highest quality output

10. Quick-Reference Prompt Library

📊 SLIDE 32–35: PROMPT LIBRARY WORKSHOP EXERCISES

Four slides, one per exercise group: (1) Alt-text generation exercise: participants photograph an object and generate alt text using Prompt 1. (2) Easy Read exercise: participants convert a library policy using Prompt 2. (3) Social Story exercise: participants create a library visit story using Prompt 4. (4) Multilingual signage exercise: participants generate signage using Prompt 5. Each slide shows the prompt template on the left and expected output format on the right. QR code linking to a shared document with all 10 prompts.

Prompt 1: Alt-Text Generation

```
Analyze this image and generate alt-text following WCAG 2.2 guidelines.

Provide THREE versions:
1. Short alt-text (under 125 characters): Focus on the image's purpose and essential
content, not just visual appearance.
2. Long description (1–3 sentences): Include context, visible text, spatial
relationships, and relevant mood/tone.
3. Extended description (for complex images): Enable a person who cannot see the image
to fully understand all information it conveys.

CONTEXT: This image will be used on a [public library website / event flyer / digital
exhibit]. The target audience is [library patrons / community members].

GUIDELINES:
- Describe PURPOSE, not just appearance
- Transcribe any visible text
- Avoid "Image of..." or "Picture of..."
- Do not describe perceived race or disability unless directly relevant
- If decorative only, state: "Recommend null alt text (alt='')"
```

Prompt 2: Easy Read / Plain Language Conversion

```
Convert the following text to Easy Read format following Inclusion Europe guidelines:
- Maximum 15 words per sentence, one idea per sentence
- Use simple everyday words; explain difficult words in brackets
- Use active voice and present tense
- Suggest a simple image after each paragraph: [IMAGE: description]
- Never use abbreviations, metaphors, jargon, italics, or ALL CAPS
- Left-align all text
- Begin with a 1–2 sentence summary of the key message
- Target reading level: Grade 3–5

Text to convert: [PASTE LIBRARY TEXT]
```

Prompt 3: Document Accessibility Audit

```
Perform an accessibility audit of this document against WCAG 2.2 and plain language
best practices. Rate each area as Pass, Needs Improvement, or Fail:

1. Reading level (appropriate for general public?)
2. Plain language (jargon avoided or explained?)
3. Heading structure (proper hierarchy for screen readers?)
4. Link text (meaningful, not "click here"?)
5. Color/formatting reliance (meaning conveyed without color alone?)
6. Abbreviations (spelled out on first use?)
7. Inclusive language (respectful disability language?)
8. Multiple contact methods provided?
```

```
Provide: summary score, detailed findings, specific rewritten examples, and priority
ranking (Critical / Important / Nice-to-have).

Document: [PASTE TEXT]
```

Prompt 4: Social Story for a Library Visit

```
Create a Social Story about visiting [LIBRARY NAME] for a [AGE]-year-old on the autism
spectrum. Follow Carol Gray's guidelines:
- Write in first person ("I will..." / "I can...")
- Never use "you"
- Ratio: 2–5 descriptive sentences per 1 coaching sentence
- Include sections: Getting Ready, Arriving, Inside the Library, Doing My Activity, If
I Need Help, If I Feel Overwhelmed, Leaving
- Include sensory descriptions (sounds, lights, textures)
- Include coping strategies for overwhelming moments
- Add [PHOTO SUGGESTION] tags for accompanying visuals
- Keep each section to 3–5 sentences
- Avoid idioms, sarcasm, or ambiguous language
- End with positive affirmation

Sensory features: [fluorescent lighting / quiet zones / automatic doors]
Available accommodations: [noise-canceling headphones / sensory kits / quiet room]
Visit purpose: [borrowing books / attending storytime / using computers]
```

Prompt 5: Multilingual Library Signage

```
Create accessible multilingual library signage for:
Sign purpose: [Wayfinding / Hours / Rules / Program Announcement]
English text: [PASTE TEXT]
Target languages: [list languages needed]

For each language provide:
- Translated text (simple, direct phrasing — no idioms)
- Text directionality (LTR or RTL)
- Pronunciation guide for staff
- Cultural considerations or connotation warnings
- Recommended minimum font size for readability

Also suggest a universal icon/pictogram description for the sign.
```

Prompt 6: Audio Description Script

```
Write an audio description script for [library video / virtual tour / exhibit]
following WCAG 1.2.3/1.2.5 guidelines.

Scene breakdown:
- Scene 1 [timestamp]: [visual content]
- Scene 2 [timestamp]: [visual content]

Rules:
- Describe ONLY what is visually present — do not interpret
- Fit descriptions into natural pauses in dialogue
- Prioritize: WHO (by role/clothing) → WHERE → WHAT → HOW
- Use present tense
- Note on-screen text that needs reading aloud
- Mark [EXTENDED DESCRIPTION] where pausing video is necessary

Output as two-column script: Timestamp | Audio Description Text
```

Prompt 7: Sensory-Friendly Library Guide

```
Create a sensory-friendly guide for [LIBRARY NAME] for patrons with sensory processing
differences. Include:

1. Sensory map: Room-by-room description of sound levels, lighting, crowds, smells,
floors
2. Best times to visit (quietest hours/days)
3. Available accommodations and how to access them
4. Coping strategies for sensory overload
5. Quiet zone locations and rules
6. Staff support (who to ask, what to say)
7. Quick exit routes if overwhelmed
8. Sensory-friendly program schedule

Library details: [floors, areas, sensory challenges, lighting type, accommodations]
Tone: Welcoming, non-clinical. Format with clear headings and short paragraphs.
```

Prompt 8: Wayfinding for Multiple Disability Types

```
Create wayfinding instructions from [START] to [DESTINATION] at [LIBRARY NAME].
Provide FOUR versions:

1. Wheelchair/mobility users: accessible entrance, elevators, door widths, floor
surfaces, ramps, restrooms
2. Blind/low vision: tactile landmarks, clock-face directions ("door at your 2
o'clock"), step counts, hazard warnings, audio cues
3. Deaf/hard of hearing: visual landmarks, color-coded signage, staff identification,
communication options
4. Cognitive accessibility: numbered steps (max 10), simple vocabulary, "Look for..."
landmarks, [PHOTO: description], reassurance language
```

Prompt 9: Catalog Record Accessibility Enhancement

```
Enhance this catalog record with accessibility metadata following Schema.org and ONIX
standards:

Title: [TITLE] | Format: [format] | Year: [YEAR]

Add: accessMode, accessModeSufficient, accessibilityFeature, accessibilityHazard,
accessibilitySummary.

Also generate: Patron-friendly "Accessibility Notes" for catalog display and MARC 21
field 532 suggestion.
```

Prompt 10: Accessible Event Announcement (Multi-Format)

```
Rewrite this event for full accessibility. Generate ALL versions:

Event: [NAME] | Date: [DATE] | Location: [LOCATION] | Cost: [COST]
Description: [PASTE DESCRIPTION]

1. Web version: H1/H2 headings, plain language, meaningful links, accommodation
statement
2. Social media: Under 280 chars + alt-text suggestion
3. Email: "At a Glance" box + full description + accessibility section
4. Print flyer: Large print (14pt+), high contrast, QR code
5. Easy Read: Grade 5, one idea per line, [IMAGE: description] tags
```

11. Ethical Foundations for AI-Powered Accessibility

📊 SLIDE 36: ETHICAL FRAMEWORK

Four ethical pillars displayed as columns: (1) Bias & Automated Ableism — NLP models score disability statements as more negative/toxic (Penn State, 2023 ACL). (2) Nothing About Us Without Us — advisory boards, participatory design, disabled co-researchers. (3) Avoiding Inspiration Porn — Stella Young (2012), medical vs. social model check. (4) Global South Considerations — only 3% access to assistive tech in LMICs, 40% of websites meet critical blind-user accessibility. Closing: "When should AI NOT be used?"

11.1 Bias, Ableism, and the Medical Model Trap

Penn State researchers (2023 ACL Workshop) found that NLP sentiment and toxicity models systematically score disability-related statements as more negative and toxic than control categories — a phenomenon termed "automated ableism." AI systems often define disability in historical medical terms while modern disability rights frameworks have moved to social and rights-based models. A fundamental paradox exists: if you do not incorporate disability data, algorithms discriminate; if you incorporate it, many people are missed because of how datasets divide on identity axes.

11.2 Nothing About Us Without Us

The phrase, popularized by disability rights activists in the 1980s–90s, demands that disabled people be partners in decisions affecting them. For library AI deployment, this means establishing advisory boards composed of representative persons with disabilities, using participatory design approaches where disabled users are co-researchers, and testing all AI-generated accessibility content with members of the target community before publication.

11.3 Avoiding Inspiration Porn

Australian disability activist Stella Young coined "inspiration porn" in 2012 to describe the objectification of disabled people for non-disabled people's benefit. AI trained on web data containing massive amounts of inspiration porn may default to these tropes. Librarians must review all AI-generated disability content for medical model framing versus social model framing, and check whether non-disabled people are positioned as saviors.

11.4 Global South Considerations

AI tools designed for Western contexts may be inappropriate for Global South settings. Only 3% of those needing assistive technology in low- and middle-income countries have access. Only 40% of evaluated websites in the Global South meet critical accessibility guidelines for blind users (arXiv, 2025). Appropriate approaches include prioritizing open-source and offline-capable tools, mobile-first solutions, and leveraging human expertise where it may be more cost-effective and culturally appropriate than AI.

11.5 When AI Should Not Be Used for Accessibility

Based on the research synthesized in this guide, AI should not be the sole authority for: legal compliance determinations; disability identification or service eligibility; mental health or crisis

support; cultural or spiritual disability content for communities where training data is scarce; consent collection from people with intellectual disabilities; high-stakes accommodation decisions; or any context where no human review is possible for the output.

12. Universal Design for Learning and Library AI Services

📊 SLIDE 37: UDL + AI IN LIBRARIES

Three-column UDL framework (CAST 3.0, 2024): Engagement (“Why”) → AI-personalized recommendations, adaptive chatbots. Representation (“What”) → AI text leveling, audio versions, translations, image descriptions. Action & Expression (“How”) → Voice-to-text search, multiple feedback modalities, AI-assisted forms. Key stat: 70% of educators want to learn AI for disability support (CAST, 2025). Principle: “Co-design from the start — do not retrofit.”

The CAST UDL Guidelines (Version 3.0, 2024) provide a design framework built on three brain networks that maps directly to AI-powered library services. Multiple Means of Engagement (the “why”): AI-personalized reading recommendations, adaptive chatbots. Multiple Means of Representation (the “what”): AI text leveling, audio versions, translations, image descriptions. Multiple Means of Action and Expression (the “how”): voice-to-text search, multiple feedback modalities, AI-assisted form filling.

CAST’s March 2025 guidance emphasizes that 70% of educators specifically want to learn how to use AI to support students with disabilities and English language learners. Their key principle: co-design from the start — do not retrofit. Include learner voices, especially people with disabilities and multilingual learners, in AI tool design from the beginning.

13. Conclusion: From Tools to Transformation

📊 SLIDE 38: CLOSING SLIDE: THREE KEY INSIGHTS

Three insight cards: (1) LOCAL AI HAS REACHED A PRACTICAL THRESHOLD — a US$600 desktop running Ollama performs text simplification, translation, and summarization offline with complete privacy. (2) THE LANDSCAPE IS ASYMMETRIC — visual accessibility has the most mature tools; cognitive accessibility (Easy Read, social stories) delivers the greatest immediate impact; sign language remains the hardest frontier. (3) THE GLOBAL SOUTH IS BOTH GREATEST NEED AND GREATEST INNOVATION SITE — Terp 360 (Kenya), Tiflonexos (Argentina), NDLI (India). Final message: "PAPAYAS reminds us every output is a simulation. The Marrakesh Treaty provides the mandate. These tools provide the means. What remains is YOUR commitment." Contact info and QR code to IAparaBibliotecas.com.

This guide provides a complete toolkit, but tools alone change nothing. The transformation happens when librarians combine critical understanding of AI limitations (PAPAYAS), practical prompt engineering skills, ethical commitment to disability rights, and genuine partnership with disabled communities.

Three insights emerge from this research that go beyond the expected. First, local AI has reached a practical threshold for libraries. A US$600 refurbished desktop running Ollama with Llama 3.1 can perform text simplification, translation, and summarization entirely offline, with complete patron privacy. This is not future technology — it is deployable today.

Second, the accessibility tools landscape is asymmetric. Visual accessibility has the most mature AI solutions (Be My Eyes, Seeing AI, Lookout), while auditory accessibility through sign language remains the most challenging frontier, with Deaf community acceptance as the critical barrier rather than technology. Cognitive accessibility — Easy Read, social stories, visual schedules — may be where AI chatbots deliver the greatest immediate impact, because structured prompts reliably produce usable first drafts.

Third, the Global South is both the greatest need and the greatest site of innovation. Terp 360 from Kenya, Tiflonexos from Argentina, NDLI from India, and community fieldworker models from Ecuador represent accessibility solutions born from constraint, which AI should amplify rather than replace.

The PAPAYAS framework reminds us that every AI output is a statistical simulation, not expert knowledge. The Marrakesh Treaty and CRPD provide the legal mandate. The tools documented here provide the means. What remains is the librarian's commitment to deploy them with rigor, humility, and authentic partnership with the communities they serve.

santiago@medejean.co