How Technology is Changing Clinical Research

Scott A. Streibich, MBA, MHS, CHC, CHRC
Director Clinical Research Revenue Cycle
Office of Clinical Trials
Johns Hopkins University School of Medicine

Objectives

- Define Generative AI
- What Risks Should Users Consider
- Practical Use Cases for Coordinators and Study Teams
- What Technology is Coming Next

What is Generative AI?

- A Prediction Engine: Think of it as "Autocomplete on Steroids." It doesn't "know" facts; it uses math to predict the most likely next word in a sentence based on patterns it learned from millions of books and articles.
- Creating vs. Searching: Unlike Google, which finds a copy of an answer that already exists, Generative AI builds a brand-new response from scratch every time you ask it a question.
- The Probability Gap: Because it is guessing the next word based on odds, it can sound very confident even when it is wrong. This is why human oversight is required—it's a powerful calculator for words, not a source of truth.

Poll Question #1

How would you describe your current familiarity with and application of Generative AI tools (e.g., ChatGPT, Gemini, Grok, Copilot, Claude) in professional or personal tasks?

1. Regular User: I have professional access and use the tools routinely in my work or research processes
2. Exploratory User: I have tried a few tools for specific tasks and understand the basic concepts ("prompting")
3. Aware Observer: I know what Generative AI is, but I have not yet used the tools myself
4. New to the Concept: I am hearing the term often, however the technology and its potential uses at work are largely unfamiliar to me

Poll Question #2

Within your current professional environment, what is your level of access to Generative AI tools?

1. Restricted/None: Access is strictly prohibited or unavailable on all work devices and networks.

2. Public Tools Only: I can use public tools (e.g., free versions of ChatGPT, Copilot, Gemini) but only for general, non-sensitive work tasks.

3. Secure, Internal Tools: I have access to institution-specific, secure Generative AI systems ("Walled Gardens") designed to use our private/internal research data and may include use of PHI.

Walled Garden AI Environments
in Academic Healthcare Organizations

What is a Walled Garden AI Environment?

- A closed, controlled AI ecosystem within an organization
- Access is limited to approved users, data and tools
- Designed for security, compliance and governance

Key Features

- Internal hosting of AI models and data
- Strict access controls and audit trails

Why Use It in Academic Healthcare?

- Protect PHI & sensitive research data (HIPAA compliance)
- Prevent data leakage to external models er vendors
- Ensure reproducibility & transparency in research
- Enable safe innovation with curated datasets and vetted models

Benefits

- Trustworthy AI for patient care and research
- Reduced risk of regulatory violations

Access to Generative AI Resources

- The Infrastructure Gap:

  - Academic Centers & Affiliate Networks : Secure 'Walled Gardens' (Enterprise AI).

  - Non-Academic Affiliates & Community Sites: Often rely on public, free tools (Risk Exposure).

- The Rule: If you or your organization aren't paying for it, YOU (and your data) are the product.

Risk #1: Hallucinations (When AI Lies)

- Hallucination: AI generating false information with high confidence.
- The Danger: Inventing inclusion criteria, fake citations, or wrong side effect %.
- The CRC Role: You are the fact-checker. Never copy-paste without verification.
- Use pre- or post- prompt validation strategies.

Risk #2: Introduces Bias

- Generative AI Models are trained on historical internet/medical data.
- The Consequence:

  - May overlook symptoms presenting differently in diverse groups.

  - Recruitment strategies may default to historical biases.

Risk #3: The Privacy Trap (HIPAA)

- The Golden Rule: NEVER put PHI into a public AI tool. Consult your walled garden policies first.
- Avoid The (18) Eighteen Protected Health Information (PHI) HIPAA Identifiers.
- The 'Black Box': Once data enters a public model, we don't know where it goes or who has access it.

Risk #4 Automation Complacency

- The Trap of Reliability: As LLMs become more accurate, users naturally stop double-checking the work, allowing subtle but high-stakes "hallucinations" to slip through unnoticed.
- The Erosion of Expertise: Over-reliance on AI for critical thinking can lead to "skill atrophy," where humans lose the deep domain knowledge required to supervise and course-correct the machines.
- From Creator to Passive Monitor: The greatest risk is shifting from an active participant to a passive observer, where the speed of AI output outpaces our ability to maintain accountability and oversight.

Use Case #1: Improving Readability (ICFs)

- Problem: Consent forms are often too dense for patients.
- Prompt: 'Rewrite this paragraph at a 6th-grade reading level. Keep medical meaning but explain terms.'
- Benefit: Better patient understanding and retention.

Use Case #2: Meeting Recaps

- Tools: Teams Premium, Zoom AI, Otter.ai (Check site policy!). AI scribes are similar applications.
- Application: Auto-generate minutes for SIVs or Monitor Visits.
- The Win: Focus on the conversation, not the note-taking. Consider the risks of private use.

Use Case #3: Training Materials

- Scenario: New protocol amendment requires staff training.
- Prompt: 'Create a 5-question multiple-choice quiz based on this protocol summary to test exclusion criteria.'
- Benefit: Rapid creation of training documentation.

Use Case #4: Operational Emails

- Scenario: Drafting difficult emails to sponsors or recruitment letters.
- Prompt: 'Draft a professional email to a sponsor explaining we need a budget amendment due to an unanticipated increase in screen failures.'
- Benefits: Overcoming Writer's Block. Improve Tone and Clarity.

Use Case #5: Document Comparison

- The Challenge: Finding subtle changes between Protocol v1 and v2.
- The Fix: Upload (redacted) versions and ask 'List changes to the Schedule of Events'.
- Tool: Most tools can accomplish this.
- Risk: De-Identifying information outside of a walled garden is impractical and time-consuming. Use this strategy in walled gardens for proprietary documents.

What’s Next: Retrieval Automated Generation (RAG)

- Retrieval-Augmented Generation is an organizational construct that allows an AI to look up external facts in real-time before generating an answer.
  - Standard AI models have a "Knowledge Cutoff" (they don't know today’s news).
  - Prone to "Hallucinations" (making up facts when they don't know the answer).
  - Improves on walled garden security with focus on internal resources
- How it Works (The 3-Step Loop):
  - Retrieve: The system searches your private documents or live data for the answer.
  - Augment: It adds that specific info to your question.
  - Generate: The AI writes a response based only on the retrieved facts.

RAG- From Search Engine to Strategic Partner

- Instant Expert Knowledge: In the future, you won’t "search" for a policy, protocol or regulatory binder
    - You will ask, and the AI will synthesize the answer instantly from internal resources.
- Hyper-Personalization: RAG will allow AI to understand your unique history, preferences, and past projects to provide context-aware suggestions that feel like they came from a long-term colleague.
- Verifiable Trust: Every answer will come with Automatic Citations. You’ll never have to guess if the AI is "lying"—you can click a link to the exact document it used for the answer.

Closing Tip

Pre- and Post- Prompt Validation Strategies

- Chain-of-Verification Prompt
- “Pre-Computation” Prompt

Chain-of-Verification

“Please review your last response for accuracy. Perform this quick check:

1. Audit: Note any names, dates, or technical details you included.
2. Verify: Confirm whether each point is based on reliable information or an assumption. If you’re unsure, say “I’m not certain about X.”
3. Support: If possible, explain the source or reasoning behind your answer.

If you find any errors, please provide a corrected version.”

Pre-Computation Check

“Before you answer, walk through your reasoning step by step.

If any part depends on facts you’re not completely sure about—or would normally need a RAG-style search you can’t do— let me know your confidence level (Very Confident, Confident, Not Confident) for each piece.”

Q&A, Contact Information & Thank You

Scott A. Streibich, MBA, MHS, CHC, CHRC

Director, Clinical Research Revenue Cycle

Office of Clinical Trials

Johns Hopkins University School of Medicine

Email: sstreib1@jhmi.edu

Linkedin: https://www.linkedin.com/in/scott-streibich/