

Three Habits for Reliable Results

AI Essentials for the Busy Clinician · Brian Whittaker · RRM Academy

Orientation

Three habits separate the people who get real value from the people who quit in a week. Here they are.

Lesson

Habit 1: Give it the details

The AI knows nothing about your situation except what you tell it (your **context**). Compare:

"Write a letter about a missed appointment."

"Write a warm, brief letter from a small women's health practice to a patient who missed her follow-up. Goal: rebook without guilt. Fifth-grade reading level, under 120 words, end with our scheduling line."

Ten seconds of detail, a completely different letter. The template: **who for, what for, tone, length, format**. Five slots.

Term: Context. Everything you have put into the conversation. The AI's entire knowledge of your situation.

Habit 2: Give it the document

The biggest upgrade in this course: **if the answer should come from a document, put the document in**. Paste it or upload it, then ask.

From its memory: decent odds of drift. From your document: reliable. This is called **grounding**, and it is your default for anything factual.

- "Summarize this paper's findings and limitations in plain language."
- "Turn this policy into a one-page staff checklist."
- "Here are our post-procedure instructions. Rewrite at a sixth-grade level. Keep every warning."

Term: Grounding. Giving the AI the source and asking it to work from that. Your most reliable results live here.

Habit 3: One chat per task

The AI's attention (**working memory**) is finite. One giant everything-chat degrades: instructions fade, documents blur together. Not a malfunction; overflow.

- New task, new chat.
- Long chat? Ask for a summary of what is settled, carry it into a fresh one.
- New chats remember nothing from old ones. Provide what it needs again.

Term: Working memory (context window). The finite span of conversation the AI can attend to. One chat per task keeps it sharp.

The cheat sheet

Five-slot prompt. Document goes in first. One chat per task. Refine, don't retype: "shorter," "warmer," "as a table."

Scenario: try it yourself

Setup. 15 minutes, three tasks: simplify a consent explainer, draft an MA job posting, pull key recommendations from a 30-page guideline PDF.

Model answer. Three chats. (1) Paste the explainer: "sixth-grade level, warm, keep every risk statement, under 300 words." (2) "Job posting, part-time MA, small women's health practice, friendly professional, duties + hours placeholder + how to apply, under 250 words." (3) Upload the PDF: "key recommendations, plain language, grouped by topic, with page numbers." Review each, refine once.

Terms from this module

Term	What it means, and why you care
Context	Everything you have put in the conversation; all the AI knows of your situation. <i>Fill it with the five-slot prompt.</i>
Grounding	

Term	What it means, and why you care
	Give it the document, ask it to work from that. <i>The default for anything factual.</i>
Working memory (context window)	The finite span it can attend to. <i>One chat per task; decay is overflow, not malfunction.</i>

From the AI Essentials course for RRM Academy members · rrmacademy.org/courses/ai-essentials · © 2026 RRM Academy.
For personal use by course participants.