

How to Master *Learning with AI*

Dr. Mark Keith, *A Student’s Complete Guide to Thinking Better with Artificial Intelligence* (v1.2)

WHAT THIS COVERS	TIME NEEDED
Sections 1, 2, 8, 9 + Appendix A	~90 minutes total, in two sittings (~45 min each)

The quiz opens before the reading is technically due. That is deliberate. Do the reading first, take the quiz, then come to lecture with the framework already in your head.

The guide is 45 pages. **You are not expected to master all of it.** Four parts carry the ideas you will be quizzed on and that we will use for the rest of the semester. The rest is a toolkit — genuinely useful, but you will reach for it later, when you need it.

WHAT TO READ CLOSELY, WHAT TO SKIM

SECTIONS	HOW TO HANDLE IT
1, 2, 8, 9 + Appendix A	Read closely. This is the quizzable core: why AI use backfires, the eight engagement modes, the self-honesty tests, and the integrity line.
3, 4, 5, 6, 7	Skim the headings. Setup, entry points, the prompt library, and cross-tool verification. This is the material the AI marketing assignment puts into your hands, so you will work through it there. Know what lives where; do not memorize it now.
10, App. B–D	Reference only. The 59-skill toolkit and the course-setup templates. Bookmark them.

SESSION ONE · ABOUT 45 MINUTES

The problem, and the model that fixes it

STEP 1 · 10 MIN · READ SECTION 1

Read Sections 1.1 through 1.4. Keith is making one argument: the tool is not the variable, *how you use it* is. Watch for the BCG consultant study and the two groups that got opposite results from the same

AI.

STOP AND ANSWER, NO PEEKING

What is “the illusion of understanding,” and what is the one-sentence test for whether you have it?

STEP 2 · 20 MIN · READ SECTION 2 + APPENDIX A

This is the heart of the course reading. Section 2 lays out the AI Engagement Model (AIEM): five dispositions at the roots, six principles in the trunk, eight modes as the branches, outcomes as the fruit, and the *Becoming* loop running through it.

Read 2.5 twice. The modes are **a web, not a ladder** — there is no “best” mode, and Mode 8 is not the top of a staircase. Note the dependencies: you cannot verify (Mode 5) material you do not yet understand (Mode 3), and you cannot stress-test an argument (Mode 7) you have not yet built (Mode 4).

STOP AND ANSWER, NO PEEKING

Close the guide. On a blank sheet, write all eight modes in order with a one-line example prompt for each, and mark which grouping each falls into. Then check yourself against Appendix A. Whatever you missed is what you review at the start of the next session.

STEP 3 · 15 MIN · DO IT, DO NOT JUST READ IT

Copy the five-part prompting brief from Section 2.7 into any AI assistant and fill in the brackets for something you are actually working on this week — a MKTG 201 concept you are shaky on works well. Name the mode you want on purpose.

Then look at your own last three AI conversations, from any class. Which mode was each one? Be honest. Most students find all three were Mode 1 or 2, which is exactly Keith’s point.

SESSION TWO · ABOUT 45 MINUTES

Self-honesty, and the line you do not cross

STEP 1 · 5 MIN · WARM UP COLD

Before opening the guide: name the eight modes out loud. Then name the three behavioral groupings and which modes belong to each. This takes two minutes and it is the difference between recognizing the table and knowing it.

STEP 2 · 15 MIN · READ SECTION 8

Three tests for real understanding, five warning signs of sliding into passivity, and the confident-tone trap in writing. Learn the three tests by name — you will use them on yourself for the rest of the semester.

STOP AND ANSWER, NO PEEKING

List the five warning signs that you have slid into passive consumption. Which one describes you most often?

STEP 3 · 15 MIN · READ SECTION 9

Integrity and disclosure. Keith draws the line in one sentence — know it well enough to quote the substance. Then pay close attention to the citation-fabrication numbers. They are concrete, they are quizzable, and they are the reason “verify before you rely” is not a slogan.

STOP AND ANSWER, NO PEEKING

Where exactly is the line between AI as a learning tool and academic dishonesty? And what makes a disclosure statement good rather than useless?

STEP 4 · 10 MIN · CLOSE THE CHAT

Shut the document. Explain the AI Engagement Model out loud to a roommate, a sibling, or an empty room, in under three minutes: what the four layers are, what the eight modes are for, and why a “better” mode does not exist. If you stall, you found your gap. Go back to that section only.

THE ONE TABLE TO KNOW COLD

If you learn nothing else from the reading, learn this. Cover the right two columns and work down the list.

#	MODE	SOUNDS LIKE	GROUPING
1	Oracle	“What’s the answer to this question?”	Passivity
2	Production Assistant	“Write this essay / summarize this chapter.”	Passivity
3	Tutor	“Help me understand this. Don’t just tell me — ask me questions.”	Partnership
4	Collaborative Problem-Solver	“Here’s my thinking so far. What am I missing?”	Partnership
5	Verification Agent	“I think this is right, but check my work.”	Agency

#	MODE	SOUNDS LIKE	GROUPING
6	Creative Expander	<i>“Five approaches I haven’t considered?”</i>	Agency
7	Critical Challenger	<i>“Attack my argument. Where is it weakest?”</i>	Agency
8	Problem Setter	<i>“Am I even asking the right question?”</i>	Agency

MASTERY CHECK — ANSWER THESE BEFORE YOU OPEN THE QUIZ

Document closed. If you can answer all twelve out loud, you are ready.

1. In the BCG consultant study, what happened to the structured group and what happened to the unstructured group? Why is the second number the more important one?
2. Why does a clear AI explanation *feel* like learning when it often is not? Name the psychological phenomenon.
3. What is cognitive offloading, and how does Keith say AI differs from the Google Effect?
4. What single learning strategy does the research identify as most effective — and why is ordinary AI use its exact opposite?
5. Name the four layers of the AI Engagement Model from roots to fruit, and say which direction the causal chain runs.
6. List all eight modes with their groupings. (See the table above.)
7. Why are the modes described as a web rather than a ladder? Give the two dependency examples Keith uses.
8. What is automation bias, and why are students especially vulnerable to it?
9. Name the five parts of the AIEM prompting brief from Section 2.7.
10. Name the three tests for genuine understanding from Section 8.1.
11. Where is the line between AI as a learning tool and academic dishonesty?
12. Roughly what share of AI-generated citations are fabricated by current models — and what share of the *real* ones still contain errors?

Scoring yourself honestly. “I remember reading that” is not an answer. If you cannot produce it from memory, you have recognition, not understanding — which is the exact failure the guide spends 45 pages warning you about. Go back to the one section that gap came from.

A note on how to use AI on this reading. You may absolutely use AI to help you master this document — that is the point. But notice which mode you are in while you do it. Asking Claude to summarize the guide is Mode 2, and it will leave you with nothing when the quiz opens. Asking Claude to quiz you on the eight modes, then to challenge your explanation of why they are a web and not a ladder, is Modes 3, 5, and 7. Same tool, ninety minutes either way, opposite outcomes. Keith opened the guide with exactly that finding.