Slow AI: a 60-minute workshop

A complete plan for one session with students. It runs with staff too, and the changes are at the foot. You need the specimen text on page 2 and nothing else. Read chapters 2 and 7 first, about 25 minutes.

Slow AI: Knowing When to Use AI and When to Leave It Alone. Sam Illingworth, 2026. Paperback, 99 pages, ISBN 9798186576105. Reviewed by Professor T. D. Wilson in *Information Research* 31(3), 548 to 549 (2026), doi.org/10.47989/ir31367985.

Learning outcomes

By the end of the session students will be able to:

1. apply four named tests to decide whether a task suits AI
2. analyse machine written text for what it asserts, assumes, and omits
3. justify, in writing, one task they will keep for themselves

Paste these into your own paperwork, changing the verbs to your level. They are Bloom's revised.

The four tests

The spine of the session. Put them on the board at minute 5 and leave them there.

	The test	Ask	It fails when
1	**Bounded task**	Where does this task stop?	the ask has no edges: be my friend, do my thinking, fill my feed
2	**Human downstream**	Who qualified catches the errors, and are they inside the decision?	nobody can be named, or the named person never sees the output
3	**Honest comparison**	Better than what is actually available here, right now?	the comparison is to nothing, or to an expert nobody can reach
4	**The thinking stays yours**	After a hundred of these, are you more capable or less?	the user ends up smaller

Chapter 2 makes the case for good AI use early on purpose. Refusing AI where all four tests pass is its own failure of judgement, so this gives students a way to say yes as well as no.

The sixty minutes

Time	What happens	What you do
0 to 5	Each student writes two lines: one task they would hand to AI this week, one they would not. Silent, no discussion.	Collect nothing. They keep the paper.
5 to 15	Teach the four tests. One example each from your own discipline.	Board, then leave it up.
15 to 30	Pairs run all four tests on the specimen text. What does it assert, assume, and omit?	Circulate. Push hard on test 2: who exactly?
30 to 40	Three pairs report back. Name which test each objection came from.	This is where the tests stick.

40 to 52	Back to their own two lines from minute 0. Run the four tests on both. Has either moved?	Expect the second list to grow.
52 to 60	Each writes one sentence: the task they will stop handing over, and which test it fails. Collected on the way out.	Those sentences are your evidence for outcome 3.

Read a handful aloud at the start of the next session. The ones that change behaviour are the specific ones.

The specimen text

> *"We understand that change can feel unsettling, and we want to reassure you that our approach to AI puts people first. These tools are designed to support you, freeing you up to focus on the work that matters most. Roles will naturally evolve as we embrace these new capabilities together, and we are committed to supporting everyone through this transition."*

A model's draft of a message reassuring staff about an AI rollout. Chapter 7 reads it to the end. Three things students usually find: it never says what the roles evolve into, it never names who decided, and freeing you up assumes the freed time comes back to you.

Running it with staff

Same sixty minutes, three changes. Give them a minute alone with the specimen text before pairing, because it is about them and they will need the minute. At minute 40 the two lines are marking and planning rather than assessed work. And the closing sentence is usually a policy question, so know where you are sending it.

If you want it assessed

Ten tasks over two weeks. Before each one they write two sentences of their own answer first: what they think, and what they most need to know. They log it, then write 1,000 words on what changed. Marked on judgement against the four tests, whether sources were checked, and what a text omits. It works unassessed, which is where to start.

The book, and how to order it

Paperback £7.99 in the UK, $9.99 in the US, €9.99 across the EU. Twelve copies for a teaching team comes to £95.88. Order through whichever supplier you already use: it is in the trade supply chain on ISBN 9798186576105, so Gardners, Askews and Holts, GOBI and any other Ingram supplied wholesaler can raise it against a purchase order and invoice your department. An electronic edition for library licensing is being set up for January 2027.

> **Want to read it first?** Ask me for an inspection copy.

Challenged on any of it? The sources are in the book. The detection evidence is Liang et al., *Patterns* 4(7), 2023. The deskilling evidence is Bastani et al., *PNAS* 122(26), 2025. · Published under a Creative Commons Attribution 4.0 International licence (CC BY 4.0): edit it, cut it, translate it, add your own logo, and keep the attribution line. · Questions, or an inspection copy: sam.illingworth@gmail.com