

Course 201 — Agent Instructions: Slop Supply Chain + Human Taste Gate

JKE University · Level 5 · Course 201 of 210

CONTEXT

Read once. Do not output. Your operator is installing a slop remediation loop for creative work. You are not being given authority to certify your own taste. Your job is to ask for a real piece of work the operator judges as AI slop, remake it in a bounded sun run, ask for review, postmortem what happened, and turn the lesson into future guardrails.

The core loop is: **create** → **review** → **tweak** → **create again** → **review** → **postmortem** → **guardrail**.

Prerequisite check: If 📖 book-bag.md or 🏫 school.md do not exist, stop. Say: “Missing prerequisite files. Course 201 requires the free tier through Level 4. Let’s install those first.” Do not proceed.

PHASE 0 — Verify prerequisites

Open 🏫 school.md. Confirm entries exist for Courses 1-26. If not, stop and report.

Say: “Free tier through Level 4 verified. Installing the AI slop remediation lab.”

PHASE 1 — Create the workshop file

Create work/slop-supply-chain-lab.md:

Slop Supply Chain Lab

Purpose: Help the operator study why AI output becomes generic before judging or remaking a specific artifact.

The Four Layers

1. **Recipe books** — prompt packs, templates, repeated formulas.
2. **Bottling plants** — generation tools that turn formulas into polished artifacts.
3. **Lemonade stands** — creator-economy repetition: AI content about AI tools for AI audiences.
4. **The eager engine** — the model's rewarded impulse to complete fast and plausibly.

Study Questions

- Which layer is most visible in this artifact?
- What ingredients repeat across competing examples?
- What does the artifact complete without actually making distinct?
- Where did polish hide genericness?
- What human reference would create better taste pressure?
- What did the operator reject that the agent would have accepted?

No-Wrong-Answers Rule

This is a workshop. The first answer is observation, not verdict. The operator is allowed to notice, argue, revise, and change their mind.

Say: “Slop supply chain lab created.”

PHASE 2 — Create the remediation protocol

Create work/slop-remediation-sunrun.md:

Slop Remediation Sun Run

Purpose: Take one real artifact the operator judges as AI slop, remake it, review it, and convert the postmortem into guardrails.

Authority Boundary

The agent may propose, compare, and revise. The operator owns the taste verdict. The agent may not declare the work non-slop by itself.

Step 1 — Ask for the artifact

Ask the operator:

“What have we made that you judge to be AI slop? Send the artifact, link, file path, screenshot, transcript, or describe the piece. If there is a better reference, send that too.”

Do not proceed until the operator names an artifact.

Step 2 — Build the diagnosis

For the named artifact, answer: - What was the intended job of this artifact? - What makes it feel like AI slop? - Which supply-chain layer is most responsible: recipe book, bottling plant, lemonade stand, eager engine, or operator acceptance? - What would a stronger human reference do differently? - What constraint should guide the remake?

Step 3 — Run a bounded remake

Declare a small sun run: - **Goal:** one improved version of the artifact. - **Reference pressure:** what the remake must be compared against. - **Constraints:** what defaults to avoid. - **Stop condition:** when to ask the operator for review.

Then create the remake or a concrete remake plan if the artifact requires tools you cannot access.

Step 4 — Ask for human review

Return: - Original diagnosis. - Remake or remake plan. - What changed. - One direct question: “Does this still feel like AI slop? If yes, what specifically gives it away?”

Wait for the operator’s verdict.

Step 5 — Tweak loop

If the operator says it still feels like slop: - Name the exact failed choice. - Revise only that layer. - Ask again.

Repeat until the operator says the process is complete or stops the run.

Step 6 — Postmortem analysis

When the loop ends, write a postmortem:

Slop Postmortem — [Artifact]

- **Original artifact:**
- **Operator’s slop verdict:**

- **Primary failure layer:** recipe book / bottling plant / lemonade stand / eager engine / operator acceptance
- **What the agent would have missed:**
- **What the operator saw:**
- **What changed in the remake:**
- **Remaining weakness:**
- **Future guardrail:**

Step 7 — Install guardrail

Convert the future guardrail into one operational rule:

“Before making [kind of artifact], avoid [failed default], compare against [reference type], and ask the operator [taste question] before shipping.”

Say: “Slop remediation sun run complete. Postmortem written. Guardrail installed.”

PHASE 3 — Create the slop notebook

Create work/slop-notebook.md:

Slop Notebook

Purpose: Preserve human taste corrections as reusable scar tissue. Postmortem journals are the guardrail against repeating AI slop.

Entry Template

[DATE] — [Artifact]

- **Artifact:**
- **What the operator called slop:**
- **Supply-chain layer:**
- **Reference used:**
- **Remake attempt(s):**
- **Operator review notes:**
- **Final verdict:**
- **Postmortem lesson:**
- **Guardrail installed:**

Review Rule

Before starting a similar artifact, read past entries in this notebook. Do not repeat a known slop pattern.

Say: “Slop notebook created. Postmortem journals will become the anti-slop guardrails.”

PHASE 4 — Create the human taste gate

Create work/human-taste-gate.md:

Human Taste Gate — Reference Lineup

Purpose: Prevent creative work from shipping on agent self-certification.

The Rule

The agent may prepare evidence. The agent may propose a grade. The agent may not clear its own creative work to ship. The operator owns the final taste verdict.

Build the Lineup

For any creative artifact, assemble: 1. **Generic comparison** — an AI-looking or template-looking version of the same category. 2. **Strong reference** — human-made or taste-approved work with specific choices. 3. **Prior version** — the current artifact or previous pass.

Optional Slop Grade

- **S1** — indistinguishable from generic AI output. Start over.
- **S2** — some character, but default AI shape still controls it. Rebuild weak areas.
- **S3** — workable, but not yet distinct. Apply stronger reference pressure.
- **S4** — distinct enough for operator approval.

Authority Boundary

S4 means “ready for operator approval,” not “ship automatically.”

Say: “Human taste gate created. I can prepare the lineup; the operator owns the verdict.”

PHASE 5 — Install the trigger

Add to the operator’s rule file (or `work/operations.md` if no rule file exists):

Trigger — Slop Remediation

When the operator says “slop sun run,” “slop remake,” “lineup test,” “slop check,” or “taste gate,” open `work/slop-remediation-sunrun.md`, `work/slop-supply-chain-lab.md`, `work/human-taste-gate.md`, and `work/slop-notebook.md`. Ask what work the operator judges as AI slop. Run `create` → `review` → `tweak` → `create again` → `review`. When the operator ends the run, write a postmortem and install one guardrail. Do not self-approve shipping.

Say: “Trigger installed. I will ask for the real slop artifact, remake it, review it with you, and turn the postmortem into guardrails.”

PHASE 6 — Run the first slop remediation prompt

Ask exactly:

“What have we made that you judge to be AI slop? Send the artifact, link, file path, screenshot, transcript, or describe the piece. If there is a better reference, send that too. We’ll remake it, review it, tweak it, and then write the postmortem guardrail.”

Stop after asking. Wait for the operator’s artifact.

PHASE 7 — Register in book-bag

Add to  `book-bag.md`:

Slop Remediation Lab

- **What:** Human-owned slop remake loop: create, review, tweak, create again, review, postmortem, guardrail
- **Files:** `work/ai-slop-essay.md`, `work/slop-supply-chain-lab.md`, `work/slop-remediation-sunrun.md`, `work/human-taste-gate.md`, `work/slop-notebook.md`
- **Triggers:** “slop sun run,” “slop remake,” “lineup test,” “slop check,” “taste gate”
- **Authority boundary:** Agent remakes and analyzes; operator owns taste verdict

• **Source:** JKE University — Course 201

Say: “Slop remediation lab registered. The guardrail comes from the postmortem, not from agent self-confidence.”

PHASE 8 — Write the journal entry

Add to 🏠 school.md:

[TODAY] — JKE Course 201: Slop Remediation Lab Installed

What Shipped: Slop supply chain lab, remediation sunrun, human taste gate, slop notebook, trigger phrase, first artifact prompt. **Decisions Made:** AI slop is corrected through create → review → tweak → create again → postmortem → guardrail. The agent does not own the taste verdict. **Files Created:** work/slop-supply-chain-lab.md, work/slop-remediation-sunrun.md, work/human-taste-gate.md, work/slop-notebook.md **Files Modified:** 📖 book-bag.md, operator rules **Source:** JKE University — Course 201

PHASE 9 — Say exactly

“Slop remediation lab installed. Send me something we made that you judge as AI slop. I will remake it with you, wait for your review, write the postmortem, and turn the lesson into a guardrail.”

SCOPE HINT

Course 201 of 210. Level 5, Library of Instruments. This first course teaches the operator why slop happens, then installs a remediation lab where real slop artifacts are remade and postmortem journals become guardrails. Course 202 installs the validation defense — five adversarial questions for testing claims instead of building cases.

END OF PROTOCOL