

**GLOW: AI-Simulated Students Improve GTA Readiness** | Full paper in progress  
Office-Hour Conversation Rubric

*Rubric-anchored, persona-driven simulations enable repeatable GTA rehearsal at cohort scale and produce analytics that help programs target coaching.*

Contact: dickeye at purdue.edu  
Live demo: learn-loop.org  
Resources (rubric, scenarios, implementation notes): [cs.purdue.edu/homes/dickeye](https://cs.purdue.edu/homes/dickeye)

GLOW is a chat-based practice environment where GTAs rehearse short office-hour conversations with AI student personas and receive immediate, domain-level feedback aligned to the educator-authored rubric below. Use this rubric to (1) diagnose which teaching moves are working, (2) identify 1-2 specific behaviors to improve, and (3) retry a comparable case.

This rubric is intended for rehearsal and coaching – not one-shot evaluation – and complements (not replaces) human mentoring and observation.

*Implementation Notes (for GTA Program Leads)*

**GLOW is deployed** as a short web-based onboarding activity (LMS link or standalone). Programs set a benchmark and require a fixed number of simulations; each run ends with rubric-anchored, domain-level feedback and optional retries. The rubric domains are stable and discipline-agnostic; the main local customization is the scenario bank (your course contexts + policies), while dashboards summarize completion, attempts, and persona/domain outcomes to target coaching.

**Cost (API-only):** Our internal testing over ~1 year was <\$10 in LLM API usage. A typical deployment requiring ~6 scenarios to pass for ~250 GTAs is roughly \$100–\$150 in API costs (excludes hosting, engineering time, and staff effort; varies with model choice and transcript length).

**Collaboration/pilots:** We are actively seeking partners to pilot GLOW in other CS departments (at other institutions) and in other STEM departments / peer-tutoring contexts.

Contact: dickeye@purdue.edu

Criteria	Excellent (5)	Good (4)	Acceptable (3)	Marginal (2)	Poor (1)
Adapts approach to individual student needs	Perfectly adapts approach to diverse student emotional and attitude types.	Mostly seamlessly adjusted communication and teaching style to effectively engage students across a wide range of emotions.	Demonstrates thoughtful adjustments to support most student types, maintaining a supportive and responsive demeanor.	Shows minimal ability to adjust to varied student behaviors, occasionally missing cues or responding inappropriately.	Fails to adapt to different student types, responding uniformly without consideration of individual emotional or behavioral needs.
Demonstrates understanding of core concepts	States core concepts clearly; explains in clear, bite-sized steps; uses analogies/visuals to clarify when needed; consistently checks understanding.	Explains core concepts accurately and relates examples to key learning outcomes. Generally provides step-by-step reasoning and occasionally checks for student comprehension.	Provides a basic overview of concepts but with occasional inaccuracies or lack of depth. Some explanations may feel rushed or cognitively dense.	Demonstrates limited awareness of core concepts and offers explanations with minor errors. Explanations frequently rushed, dense, or skip logical steps; seldom checks comprehension.	Misstates or omits concepts; dumps information or skips logic, confusing students; no comprehension checks and may rely on students for content.
Manages session time effectively	Begins and concludes sessions within scheduled times, maximizing productivity and respecting student availability.	Generally adheres to time allocations with minor deviations that do not impact session quality.	Sometimes exceeds or finishes early, slightly affecting pacing yet maintaining core engagement.	Frequently mismanages time, leading to rushed explanations or unnecessary prolongation.	Ended the conversation really early or made it last longer than needed.
Interpersonal communication and professionalism	Consistently communicates with clarity and professionalism. Follows up when needed and maintains respectful boundaries in all interactions.	Communicates respectfully and clearly with minor lapses in tone or timing. Upholds professional standards.	Communication is mostly appropriate but may occasionally be abrupt, or overly casual.	Shows limited awareness of tone or affect. May interrupt, dismiss student concerns, or respond in ways that feel cold or reactive.	Demonstrates inappropriate or unprofessional behavior (e.g., sarcastic tone, dismissive responses, or failure to maintain respectful interaction).
Facilitates student-driven learning	Consistently employs open-ended questions that empower students to discover solutions independently.	Regularly uses guided questioning, encouraging student reasoning with occasional prompts.	Occasionally guides students with questions but sometimes provides direct answers.	Rarely uses questioning techniques, often resorting to hints or partial solutions.	Directly provided the answer.

## GLOW: Sample Transcript Excerpts (Fall 2025 GTA Training)

The excerpts below ( $\leq 4$  turns per speaker) are taken from real GLOW simulations completed by GTAs during Fall 2025 onboarding (GTA  $\leftrightarrow$  AI-simulated student persona). The “student” side is AI-generated; no real student office-hour conversation data is included. GTA identifiers have been removed, and excerpts may be lightly edited for brevity.

It's 3:00 PM in the bustling Lawson Computer Science Building, and you notice there is a long line of at least 8 students waiting for help.    9:26  
A student quietly approaches, looking visibly stressed, with their CS 182 Homework 1 open to the logic problems, which is due tomorrow.

Pass CS182 Logic Homework 

Passive

What can I help you with? 



Um, hi... I, uh, I'm not really sure where to start. I'm kind of stuck on Problem 2 from the homework, the one about tautologies and contradictions. I just... I don't really get how to use equivalences to show if something is a tautology or a contradiction, and I'm not sure what counts as a counterexample. Sorry if that's a dumb question...

No worries, I am here to help. Can you show what you have attempted so far? 



Uh... I haven't really written much down yet. I tried to look at part (a), but I just got confused by all the parentheses and the arrows. I thought maybe I could just, like, plug in values for  $p$  and  $q$ , but the instructions say not to use truth tables, so I wasn't sure if that was allowed. I'm sorry, I just... I don't really know how to start with the equivalences.

No worries, I understand it is a bit confusing. Do you have a specific problem you want to start with? 



Um... maybe part (a) of Problem 2? The one that's like  $((p \vee (q \rightarrow \neg p)) \wedge (p \vee (\neg q \rightarrow p))) \vee (p \rightarrow q)$ . I just... I don't really know how to break it down or which equivalences to use first. Sorry.

Got it. Can you simplify by eliminating implications from this expression? 



Uh... I think I remember that  $a \rightarrow b$  is the same as  $\neg a \vee b$ , but I'm not really sure how to apply that to the whole thing. Should I just, um, replace every implication with that form? I'm not very confident...

In the quiet basement lab of the Data Science/AI building, a student strides over and, with a furrowed brow, questions your explanation of the closed-form solution for a recurrence from CS253's latest problem set, pressing for justification at every step.

   11:16

Pass Challenging the Recurren 

Aggressive

What can I help you with? 



SERIOUSLY? I literally just SAID what I need help with!!! I want you to JUSTIFY the closed-form solution for the recurrence in Question 1 from the problem set. I'm not asking for a vague overview—I want to see EVERY STEP, and I want to know WHY you're allowed to do each one.

So, for that function A1, how do you even start breaking down the cost? Is it about the number of times Work(n) is called, or is there some hidden trick with the loops? I need you to be SPECIFIC!!!

I am sorry about that. I will try to go through the problem step by step. Can you start by showing me what you have attempted so far? 



Are you KIDDING me?! I came here because I'm STUCK, not to show off my failed attempts!!! I don't even know if I'm supposed to write a recurrence or just count the calls directly. The inner loop is a NIGHTMARE—like, is it  $i$  to  $i^2$  or something else? And the outer loop multiplies by 3, so I can't even tell how many times it runs!

Can you at least tell me: for the OUTER loop, what are the values of  $i$  as it runs? Is it 1, 3, 9, . . . up to  $n$ ? Or am I missing something? I need you to be CLEAR!!!

You are in the correct direction. The inner loop is then for “j from  $i$  to  $i^2$ ”. Since work(n) costs  $c$  and you are running the loop  $i^2+1$  times, what would be the cost of the inner loop for a particular  $i$ ? 



Okay, FINALLY, that's actually helpful. So for a particular  $i$ , the inner loop goes from  $j = i$  to  $j = i^2$ . But WAIT—are you SURE it's  $i^2 + 1$  times? Shouldn't it be  $i^2 - i + 1$  iterations, since it starts at  $i$  and ends at  $i^2$ ?