From Compliance to Connection: A Live Workspace for AI-Personalized Assignment Design in Doctoral Education

Daniel Umana
Ed.D. student, School of Education, American University

Abstract

Many university students now report using generative AI for academic work, and the dominant institutional response, detection and policy, sits downstream of the behavior it targets (Stanford HAI, 2026). I built and deployed a different response, a free, browser-based workspace that personalizes a doctoral writing assignment around each student's own experience before the writing begins, using a large language model to design the prompt and never to write the response. The design puts Self-Determination Theory (Ryan & Deci, 2017) and culturally responsive teaching (Hammond, 2015) into practice in four steps: Frame, Listen, Rewrite, Hand off. The workspace ran live in an American University Ed.D. course in June 2026, members of my cohort opted in and used it to complete a graded positionality memo with the course instructor's prior awareness, and I used it to write my own. This paper traces the theory into the build, walks one cycle end to end, states what the instance does and does not demonstrate, and specifies the pilot study the design now makes possible.

I. The Problem Sits Upstream of Detection

Four in five university students report having used generative AI on academic work, and formal education is lagging behind (Stanford HAI, 2026). The dominant institutional response has been detection tools, revised honor codes, and policy language in syllabi, all of which sit downstream of the behavior. Meanwhile 57 percent of students say their assessments come with inadequate guidance on whether and how AI can be used (Digital Education Council, 2026), and students who avoid these tools most often cite fear of getting caught and the perception that use is cheating, deterrents that operate entirely through external threat (Golding et al., 2025).

The research on academic dishonesty locates the conditions that matter further upstream. Students cheat more when they are extrinsically motivated, when classrooms reward performance over mastery, and when tasks read as busywork (Murdock & Anderman, 2006). Their model includes moral beliefs and peer norms too, so motivation is one branch of the account rather than the whole of it, and it is the branch assignment design can reach. A generic prompt that routes through nothing a student carries produces a low return on the student's investment, and generative AI is exceptionally good at producing the surface artifact such a prompt asks for. The assessment literature has started moving the same direction. Chirikov, Smirnov, and Kizilcec (2026) argue from a survey of 95,513 students that AI misuse calls for assessment reform rather than detection; in their data, 37 percent of students reported at least monthly use, a measure of routine practice rather than the ever-used behavior behind the four-in-five figure. Perkins, Roe, and Furze (2025) offer a widely adopted scale for redesigning assessments around explicit levels of permitted AI use. Both redesign the task itself, and the work reported here continues down to the motivational conditions the task creates, in running software rather than in principle.

Chirikov and colleagues' reform recommendations are discipline-specific and often point toward supervised formats, which this design complements rather than replaces. Their data also showed higher generative AI adoption among male, White, and Asian students, with lower adoption among women and underrepresented students of color, so motivational disconnection is one driver among several and unequal access to AI literacy is an equity problem of its own.

II. The Theory, Adapted to the Build

Why a generic prompt invites offloading

Ryan and Deci (2017) describe amotivation arising when students do not believe they can succeed or see no value in the task, and empirical work in academic settings finds the same structure: deficient ability beliefs, low task value, and unappealing task characteristics operate as distinct reasons for disengagement (Legault et al., 2006). An assignment delivered as a standardized prompt can trip all three at once. When the prompt offers no way in, the task feels beyond reach. When it connects to nothing the student is working on, the value stays invisible, and when the format reads as a requirement to satisfy, delegation to a language model becomes a rational move. The memo this project scaffolds starts well ahead of that baseline, because its subject is the student's own experience and its prompt already ties the analysis to the student's later research. The design bet of this project is that a personalized entry point can extend that strength to the remaining conditions, giving each student a reachable first move and a reason to write that runs through their current inquiry.

Where the three needs live in the tool

Self-Determination Theory identifies three psychological needs that, when satisfied, let motivation organize around work: autonomy, competence, and relatedness (Ryan & Deci, 2000, 2017). The workspace addresses each need at a specific point in its flow, and the mapping is the design:

Need	Where it lives in the workspace	What it looks like
Autonomy	Intake (Step 2) and the personalized prompt (Step 3)	The entry point is the student's because the student supplied it. The prompt opens with a rationale for the memo that runs through their own inquiry, so the reason to write is theirs before the writing starts.

Need	Where it lives in the workspace	What it looks like
Competence	Frame (Step 1) and the rubric self-check (Step 4)	The three objectives are visible before any writing happens, each intake screen offers a starter line and a worked example, and the self-check ends with the fastest next win. The next move is always reachable.
Relatedness	The generated questions (Step 3) and the peer move (Step 5)	The questions quote the student's own phrasing and refer to the settings and people the student described, and the packet closes with a draft-trade prompt that sends the student to a peer before submission.

Relatedness in SDT runs through people, and content relevance by itself maps closer to value and identified regulation, so the tool claims relatedness support only at the two points where people are actually involved: questions that treat the student's own settings and communities as legitimate academic material, which is the belonging signal Walton and Cohen (2007) documented as consequential for students of color and first-generation students, and the closing peer draft-trade move. Belonging interventions hold up best where institutions give students opportunities to belong (Walton et al., 2023), so the signal an assignment sends is a contribution rather than the whole story.

Hammond, in the design

Hammond's (2015) central claim is that the brain wires new analytical work into existing schema, the background knowledge built from lived and cultural experience, and that prompts connecting to nothing a student carries impose cognitive load that competes with the intended learning. The claim rests on one of the most established findings in learning science: new learning builds on prior knowledge (National Research Council, 2000). For this specific assignment the fit is unusually clean. The memo asks students to analyze a critical racial moment from their own schooling, so the student's racial memory is the schema and the memo is the new analytical work being wired into it. The personalization makes that wiring explicit rather than accidental.

Hammond distinguishes surface culture from deep culture, the worldview and core beliefs through which meaning gets filtered, and locates culturally responsive teaching at deep culture.

A four-question intake reaches lived context, closer to the surface than to worldview, and Moll and colleagues built the funds of knowledge concept on teachers visiting households over time, a method no form replaces (Moll et al., 1992). The workspace is an entry-point intervention. It opens the assignment in a context the student recognizes, and deep cultural connection stays the work of instruction and relationship.

Negentropy, the state the design aims at

I borrow the term negentropy from Schrödinger (1944), who used it for how living systems generate internal order against the pull toward disorder, and I use it here for the motivational state both theories describe from their own directions. When the conditions of an assignment frustrate the three needs, motivational energy dissipates: students comply at the minimum or route the task to a model, and the assignment produces a document without producing learning. When the prompt routes through what the student carries, the energy organizes around the work, which Ryan and Deci (2000) describe as movement from external toward integrated regulation. In this paper negentropy refers to a property of the conditions, the part a designer can engineer, rather than the state of any one student. It is the goal the four steps below are built to serve, and it has a practical corollary for AI use: a student writing about their own seventh-grade classroom has a reason to do the work themselves, and a model that has never been in that classroom stops being a shortcut to a passing artifact.

III. The Assignment

The workspace scaffolds an assignment: Positionality Memo #1, Critical Racial Experiences, in EDU 703, Critical Race, Anti-Racism, and Anti-Bias Pedagogy, a foundational course in American University's Ed.D. in Education Policy and Leadership. The memo is one page, single-spaced, first person, worth 200 points, and asks the student to describe critical racial moments from their own schooling and analyze how those moments shape their lens toward the purpose of education and the interpretation of data in their later dissertation work. The prompt emphasizes that the reflection must focus on the culture of racism experienced, the norms, expectations, and everyday treatment, rather than on formal systems and policy.

Three fixed objectives are within the prompt, and the personalization protects them: reflexive awareness (identify a specific critical moment and the disposition it produced), culture-of-racism analysis (analyze the cultural dimension as distinct from formal systems), and lens-to-inquiry (articulate how the lens will shape interpretation in the student's own research). The rubric has three criteria at three performance levels and does not change from student to student. The fixed rubric anchors both competence and equity: every student is evaluated against the same standard, and every student can see exactly what the strongest level of each criterion demands before writing a word. The connection to the student's own schooling and later research is the assignment's design decision, and the workspace inherits that commitment rather than inventing it. Personalization moves the entry point while the standard stays fixed.

IV. The Tool

Four steps

The workspace implements the framework as four steps. Frame: pull the assignment's fixed objectives up where the student can see them, and flag the culture-versus-systems distinction at every step. Listen: collect the student's context through four short intake questions, the settings and roles that shaped their racial experience, one critical moment, their current inquiry, and the interpretive lens they think it created. Rewrite: generate a personalized version of the assignment prompt from that intake. Hand off: deliver a Word document the student writes the memo inside, with the personalized prompt, their own intake preserved for reference, the unchanged rubric, a visible writing space, and the original assignment at the back.

The build

The workspace is a Next.js 14 application in TypeScript, deployed on Vercel, live at https://edu703memo.vercel.app (Umana, 2026). It is public, free, requires no account, and stores nothing: no login, no database, no analytics. The personalization calls Claude Sonnet 4.5 server-side through an API route, and the Word document is generated client-side in the browser and downloaded directly, so the workspace itself never stores an intake. One transit remains: a student's answers, including whatever they choose to disclose about their racial experience, pass through a commercial model provider's servers once for the generation call. The architecture cannot remove that pass, and any classroom deployment should tell students about it before they type. A student's full cycle, landing to download, takes five to ten minutes. Each intake screen offers a starter line that pre-fills a half-sentence opener and a worked example from a fictional student, Daniela Flores, so no one faces a blank box. An always-visible About panel explains the theory, what the AI does, and what it never does.

Disciplining the model

Whether the tool produces personalization or a template with swapped nouns is decided in the system prompt and the validation around it, and this is where most of the design effort went. The generation instructions require between three and six reflection questions, sized to what the intake can sustain, with four as the default. They permit short verbatim quotes from the student's intake inside a question when the quote sharpens the analytical demand, and they require the questions to refer to the specific actors, settings, ages, and decisions the student described. They require every question to demand work at the rubric's strongest level on at least one criterion, and the full set to exercise all three criteria. They forbid inventing details the student did not provide, forbid SDT and Hammond vocabulary in the output so the prompt reads as plain academic prose, and enforce the culture-versus-systems distinction as a hard constraint. The API route validates the shape of every response before it reaches the student. When the model is unavailable or a call fails, a deterministic fallback assembles a structurally identical four-question prompt through string interpolation, so the entire scaffold, wizard, Word document, and rubric self-check works with no AI at all.

What the student walks away with

The download is a Word document built to read like a course handout. It opens with the personalized prompt: a short framing of why this memo exists for this student, a focus statement, the reflection questions, a description of what a strong memo does, and a before-you-submit note that sends the student to trade drafts with a peer. The student's own intake answers follow, preserved verbatim and visually set apart, then the rubric under an explicit header stating it does not change across students, then a writing space, then the original Canvas assignment. On screen, the student can also run a rubric self-check, rating themselves on each criterion and receiving the fastest next win, and can copy a brainstorm prompt to their clipboard, pre-filled with their intake and explicit instructions to whatever AI they paste it into: ask sharp questions, never draft. The design accepts that students will use AI and gives that use a sanctioned lane pointed at thinking. The AI designs the prompt, the student writes the memo, and the tool is never called during the writing.

V. The Actual Cycle in Action

I hold three positions in what follows: author of the framework, builder of the tool, and student in the course it scaffolds. The observations come from inside the instance, and the outcome claims are withheld accordingly. In a course about positionality, I am treating that overlap as material to interrogate rather than a conflict to hide, and it maps directly onto the reflective-practitioner and culturally-responsive-research competencies my program names.

> **One cycle, traced. The intake. In my Spanish for Spanish speakers class in middle school, a student asked our teacher, a woman from Spain, whether she would join an immigration rights protest on the National Mall. She answered, "no I am not like you." In the intake I wrote that the moment made me "more perceptive to perspectives," and I described my current inquiry: how higher education institutions are approaching critical AI literacy, and whether those approaches address equity and access. What the tool generated. Four reflection questions, each built from that intake. One asked what the classroom's culture allowed the teacher to say without challenge or consequence, which is the assignment's culture-versus-systems hinge applied to my specific room. One returned my own phrase and pushed past it, requiring me to describe the interpretive reflex the moment installed instead of restating that I have one. One connected the lens to my environmental scan, asking what I would look for in policy language and resource allocation that others might miss. The last asked what my vigilance might cause me to overlook, a gap my intake never addressed. What I wrote. A one-page memo, in first person, submitted to Canvas for a grade, that answers those four questions and connects the lens they surfaced to how I will read institutional AI literacy policies. The AI designed the prompt and never saw the memo I wrote from it.**

The second question did not accept my intake's conclusion. It returned my own phrase and required the mechanism underneath it, which is the rubric's strongest level of reflexive awareness applied to me specifically. The fourth asked what my vigilance costs me, a question my intake never raised and a fill-in template could not have produced, and that difference is what the generation constraints buy.

Members of my doctoral cohort told me they used the workspace to complete the same graded memo. Their use was voluntary, and the course instructor knew about the workspace before any of them used it. Because the deployment stores nothing, no usage counts, comparisons, or outcome data exist, and this paper claims none; what I know about cohort use comes from what cohort members reported and showed me directly. A positionality memo is also a friendly first case for this framework. Positionality is by definition something only the student can supply, so the assignment and the personalization converge, and that convergence reflects the assignment's design as much as the tool's. They will not converge like that in a statistics course, and the harder, more informative test is an assignment whose content is not inherently personal.

VI. What This Settles, and What It Does Not

The build settles feasibility. The four steps run end to end in production, and repeated test intakes during development confirmed that different intakes produce distinct prompts, with different questions arising from what each intake supplied. The scaffold degrades gracefully to a no-AI fallback. Doctoral students used the workspace on a graded assignment with no requirement pushing them to.

The build settles none of the outcome questions. Whether personalized prompts change engagement, rubric performance, or AI offloading is unmeasured here. The evidence base for context personalization generally is positive but modest: experimental work found gains in situational interest, perceived value, and effort, largest for students with the lowest baseline interest, and some later studies found the effects inconsistent (Høgheim & Reber, 2015). Situational interest is an entry point, and this design treats it as exactly that. The claim under test is the combination of a reachable entry point with need-supportive structure. Two risks carry over from the framework. A language model asked to personalize can stereotype, swapping in a demographic marker and calling it relevance; the intake design reduces that risk by working only from what students volunteer about goals, work, and communities, never from inferences about identity, and the no-invented-details constraint exists for the same reason. And a self-serve tool has no per-student human review between generation and student, the gate an instructor-run deployment would need to restore.

VII. The Pilot This Design Makes Possible

The next step is specified by what is missing. The research question: does a personalized prompt, generated under these constraints, change student engagement, rubric performance, and self-reported AI use relative to the standard prompt on the same assignment? The potential design: two sections of the same course, or random assignment within one section where the instructor judges it ethical, one arm receiving the standard prompt and one the personalized version, identical rubric, blind scoring by raters who do not know the arm. Measures: rubric scores by criterion, a short basic-needs-satisfaction scale administered after submission, time-on-task where the learning platform reports it, and an anonymous post-submission self-report of how AI was used, anchored to the sanctioned brainstorm lane the tool itself provides. The intake instrument would need consent language and IRB review as human subjects research; the current tool's stores-nothing architecture means a pilot adds data collection deliberately rather than switching surveillance off. Effect sizes from the context-personalization literature suggest powering for modest effects concentrated among students with the lowest baseline interest, which is also where the equity stakes sit.

VIII. The Scaling Path

The self-serve workspace is one deployment of the framework. The institutional version is instructor-driven. In that version the instructor uploads an existing assignment, the AI breaks the assignment down into its objectives and surfaces the cultural assumptions embedded in the task context, and the instructor verifies that analysis before anything proceeds. Intake goes out through the course's own tools. The AI drafts one personalized version per student against the verified objectives, using the instructor's existing rubric, drafting one only where none exists, and the instructor reviews every version before students see it. That per-student review gate is exactly what the self-serve instance lacks, and it is the control scale depends on, both for rigor drift and for catching the stereotyping failure mode. The build reported here demonstrates that every automated step in that pipeline works today. What scale adds is the human gate and the workload question that comes with it, which the pilot's instructor-time logging should capture.

IX. Conclusion

The detection cycle treats student AI use as a discipline problem and spends institutional energy downstream of the behavior, while the design reported here treats it as a design problem and spends that energy upstream, in the prompt, where motivational conditions are set. The theory said personalization around what a student carries should reorganize the reason to do the work. The build shows the theory can run as software, cheaply, on a real assignment, with the standard held constant and the AI kept out of the writing. I wrote my own graded memo inside it, and my cohort wrote theirs. Going forward, my focus is on the pilot: putting numbers on what the design so far only makes possible.

References

Chirikov, I., Smirnov, I., & Kizilcec, R. F. (2026). Generative AI use and misuse call for assessment reform in higher education. *Science, 392*, 818–820. https://doi.org/10.1126/science.aec5115

Digital Education Council. (2026). *AI in higher education global survey 2026.* https://www.digitaleducationcouncil.com/resource-library-items/ai-in-higher-education-global-survey-2026

Golding, J. M., Lippert, A., Neuschatz, J. S., Salomon, I., & Burke, K. (2025). Generative AI and college students: Use and perceptions. *Teaching of Psychology, 52*(3), 369–380. https://doi.org/10.1177/00986283241280350

Hammond, Z. (2015). *Culturally responsive teaching and the brain: Promoting authentic engagement and rigor among culturally and linguistically diverse students*. Corwin.

Høgheim, S., & Reber, R. (2015). Supporting interest of middle school students in mathematics through context personalization and example choice. *Contemporary Educational Psychology, 42*, 17–25. https://doi.org/10.1016/j.cedpsych.2015.03.006

Legault, L., Green-Demers, I., & Pelletier, L. (2006). Why do high school students lack motivation in the classroom? Toward an understanding of academic amotivation and the role of social support. *Journal of Educational Psychology, 98*(3), 567–582. https://doi.org/10.1037/0022-0663.98.3.567

Moll, L. C., Amanti, C., Neff, D., & Gonzalez, N. (1992). Funds of knowledge for teaching: Using a qualitative approach to connect homes and classrooms. *Theory Into Practice, 31*(2), 132–141. https://doi.org/10.1080/00405849209543534

Murdock, T. B., & Anderman, E. M. (2006). Motivational perspectives on student cheating: Toward an integrated model of academic dishonesty. *Educational Psychologist, 41*(3), 129–145. https://doi.org/10.1207/s15326985ep4103_1

National Research Council. (2000). *How people learn: Brain, mind, experience, and school* (Expanded ed.). National Academies Press. https://doi.org/10.17226/9853

Perkins, M., Roe, J., & Furze, L. (2025). Reimagining the Artificial Intelligence Assessment Scale (AIAS): A refined framework for educational assessment. *Journal of University Teaching and Learning Practice, 22*(7). https://doi.org/10.53761/rrm4y757

Ryan, R. M., & Deci, E. L. (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology, 25*(1), 54–67. https://doi.org/10.1006/ceps.1999.1020

Ryan, R. M., & Deci, E. L. (2017). *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Press.

Schrödinger, E. (1944). *What is life? The physical aspect of the living cell.* Cambridge University Press.

Stanford Human-Centered AI. (2026). *Artificial intelligence index report 2026*. Stanford University. https://hai.stanford.edu/ai-index/2026-ai-index-report

Umana, D. (2026). *EDU 703 memo workspace* [Web application]. https://edu703memo.vercel.app

Walton, G. M., & Cohen, G. L. (2007). A question of belonging: Race, social fit, and achievement. *Journal of Personality and Social Psychology, 92*(1), 82–96. https://doi.org/10.1037/0022-3514.92.1.82

Walton, G. M., Murphy, M. C., Logel, C., Yeager, D. S., et al. (2023). Where and with whom does a brief social-belonging intervention promote progress in college? *Science, 380*(6644), 499–505. https://doi.org/10.1126/science.ade4420