From Compliance to Connection:

AI-Supported Assignment Design for Student Autonomy

and Cultural Relevance in Higher Education

Daniel Umana

American University

Abstract

What are students telling us when they use AI to finish their assignments? This paper argues we have been asking the wrong question about that behavior, and that the right one leads somewhere more useful than detection.

According to the Stanford AI Index 2026 Report, 4 in 5 university students now use generative AI, and over 80% of U.S. high school and college students report using it for school-related tasks (Stanford HAI, 2026). Most institutional responses have focused on detection and policy. This paper argues that detection treats the symptom. The condition underneath is motivational and cognitive, and it falls hardest on students whose academic contexts have never connected to who they are.

Drawing on Self-Determination Theory (Ryan & Deci, 2017) and Zaretta Hammond's framework connecting culturally responsive teaching to brain-based learning (Hammond, 2015), this paper proposes a generative AI-supported workflow for redesigning existing assignments so they satisfy three basic psychological needs: autonomy, competence, and relatedness. The paper extends both frameworks by introducing negentropy as a motivational goal, the generation of internal order that sustains genuine engagement without external enforcement, and argues that generative AI is the mechanism that makes this feasible at scale.

I. AI Has Already Changed How Students Approach Assignments

The Stanford AI Index 2026 Report documents that 4 in 5 university students now use generative AI, and that formal education is lagging behind (Stanford HAI, 2026). Only half of middle and high schools have AI policies in place. Just 6% of teachers say those policies are clear. The institutional response has been mostly reactive: detection tools, updated honor codes, policy language added to syllabi.

These responses treat AI use as a compliance problem. It is also a motivation problem. When a student uses AI to complete an assignment, one of two things is happening. Either they are using a tool they have access to, the way they would use any resource, to move toward an answer. Or they are offloading work they see no reason to do themselves. Both responses are rational given the motivational conditions many assignments create. Treating either as an integrity failure misses what the behavior is communicating.

Ryan and Deci (2017) identify three forms of amotivation relevant here. The first is a felt lack of competence, students may decide to not engage because they do not believe they can succeed. The second is a lack of relevance, if the task holds no connection to anything the student values or needs. The third is defiance, motivated nonaction in response to environments that thwart autonomy. All three may make AI delegation more attractive. An assignment a student cannot connect to could be an assignment they will find another way through.

Hammond (2015) adds an additional dimension not found in SDT. The brain does not process information neutrally. It searches for connection to prior experience, personal significance, and cultural schema. When an assignment draws on contexts foreign to a student's deep cultural framework, it imposes a cognitive load that competes directly with the learning the assignment is supposed to produce. Could the disengagement patterns we sometimes associate with AI use, offloading, bypassing, minimal compliance, stem at least in part from chronic cultural disconnection that long preceded the technology? Hammond's research on dependent learners suggests that students who have repeatedly encountered academic environments that do not reach their funds of knowledge develop disengagement as a learned response, not a personal failing (Hammond, 2015). Generative AI may be the newest and most efficient tool for acting on a pattern that was already there.

What conditions would make work worth doing without AI? This paper argues that generative AI, used differently, is part of the answer.

II. What Self-Determination Theory Actually Requires

SDT is not a motivational philosophy. It is an empirically grounded theory with strong claims: when autonomy, competence, and relatedness are satisfied, people engage more deeply, persist longer, and internalize the value of their work (Ryan & Deci, 2017). When they are frustrated, motivation degrades. The research on this is consistent across contexts, age groups, and cultures.

Autonomy

Autonomy in SDT does not mean freedom from structure. This distinction matters especially in educational settings, where structure is sometimes read as control. Ryan and Deci are precise on this point: the problem is not structure itself but structure that lacks rationale, that is experienced as imposed rather than reasoned. A course with clear expectations and well-explained requirements can be highly autonomy-supportive. A course that offers choices with no pedagogical grounding can feel just as controlling if students sense the choices are arbitrary. The difference lies in whether the student experiences the structure as something that serves their goals or something that serves institutional convenience.

Ryan and Deci distinguish a spectrum from external regulation, where behavior is driven by reward or threat, through introjected and identified regulation, to integrated regulation, where a person acts from their own values. The quality of learning correlates with where on that spectrum engagement sits. A student submitting an assignment to avoid failure is externally regulated. A student submitting because they see genuine connection to something they care about is moving toward integration. The difference in learning quality between those two states is not marginal.

What an assignment can do to support autonomy is offer meaningful choice, provide rationale beyond institutional authority, and acknowledge the student's own frame of reference. In the workflow described in Section VI, the instructor retains full control over the learning objectives and the evaluation criteria. That structure is preserved by design. What changes is how the student encounters it, through a context that connects to their own experience rather than one assigned without explanation.

Competence

Competence requires optimal challenge: hard enough to require genuine effort, close enough to a student's existing capacity that success feels possible. Assignments designed for an imagined average student do not offer that calibration to every student in the room. A student who finds no foothold, no existing knowledge or experience to connect to the new material, is not experiencing productive struggle. They are experiencing the first form of amotivation, the felt inability to succeed. Connecting an assignment to what a student already knows gives them a place to start. The intellectual demand stays intact. The entry point becomes accessible.

Relatedness

Relatedness is the most misread of the three needs in educational contexts. Ryan and Deci (2017) define it as the felt sense of connection to the social world one inhabits. In practice, this means learning connects to who the student is, who they are becoming, and the problems and communities they have already decided matter. Surface personalization does not produce this. Replacing a generic example with one that names a student's demographic group performs diversity without creating connection.

Relatedness also requires what Walton and Cohen (2007) call a sense of social belonging, the felt confidence that one is accepted and valued in an academic setting. Their research documents that belonging uncertainty, the chronic concern that one does not fit or is not meant to be there, disproportionately affects students of color and first-generation college students, and that it undermines motivation and academic performance even when students have the capacity to succeed (Walton & Cohen, 2011). An assignment that connects to a student's own frame of reference is not just a motivational tool. It is a signal that the student's knowledge and context are legitimate in this academic space. That signal is part of what makes relatedness feel real rather than performed.

III. What Hammond Adds

Ryan and Deci explain what students need psychologically. Hammond explains what happens neurologically when those conditions are absent, and why the absence falls hardest on culturally and linguistically diverse students.

Hammond's (2015) central argument is that culture is the vehicle through which the brain processes and stores new information. The brain is actively looking for connection to prior experience, to schema built from deep cultural norms and ways of making meaning. When new information connects to existing schema, comprehension is faster, retention is stronger, and higher-order thinking becomes accessible. When it does not connect, cognitive load increases. What looks like disengagement is often the brain working too hard on the wrong problem.

Moll and colleagues' concept of funds of knowledge is central here. Students arrive in postsecondary classrooms carrying historically accumulated, culturally developed bodies of knowledge from their homes and communities (Moll et al., 1992). Those funds are cognitive assets. An assignment that never reaches them is leaving the student's most available intellectual resources untouched.

Hammond also makes a political argument that SDT does not. The disproportionate number of culturally and linguistically diverse students who arrive at higher education without the independent learning skills postsecondary work demands is not a deficit they carry. It is a product of educational systems that consistently denied them opportunities for higher-order thinking and productive struggle (Hammond, 2015). I work with these students. I see this regularly. They are not underprepared because they lack capacity. They are underprepared because the system undertaught them, and the assignments they now encounter in college were not designed with their histories in mind.

Where SDT and Hammond converge is the relatedness need. Hammond's argument is that cultural relevance is the neurological mechanism through which relatedness operates. An assignment that connects to a student's funds of knowledge and deep cultural schema satisfies relatedness in a way that feels real, because it is. The cognitive work connects to something the student already carries.

IV. Negentropy as the Goal

Both frameworks point toward a motivational state that neither of them name directly. I am calling it negentropy, borrowed from systems theory, where it describes the generation of internal order within a system. Applied to student motivation, entropy is what happens when the conditions of learning frustrate basic needs and cultural relevance. Motivational energy dissipates. Students comply at the minimum, or stop complying altogether. The assignment produces movement without learning.

Negentropy is the opposite state. When autonomy, competence, and relatedness are satisfied, and when learning connects to a student's own cultural schema and funds of knowledge, motivational energy organizes around the work. Ryan and Deci (2017) describe this movement as the shift from external to integrated regulation. In theory, the student may no longer do the work because the syllabus demands it. They are doing it because it connects to something they have already decided matters. That is the internal order this framework is designed to produce.

The distinction has direct implications for AI use. A student in an integrated motivational state has a reason to do the intellectual work themselves. A student externally regulated by grade threat has no such reason, and if a faster path to a passable submission exists, they will take it. Negentropy is not a soft or aspirational outcome. It is the condition that determines whether an assignment produces learning or produces a document.

The goal of the framework proposed here is to engineer the conditions for negentropy: assignments that give students a genuine reason to organize their effort around the intellectual work because the work connects to something they already carry. The learning objectives stay intact. The pathway to them changes.

V. Why Generative AI Makes This Feasible

The argument for SDT-aligned, culturally responsive assignment design is not new. The barrier has always been capacity. An instructor cannot realistically redesign 30 assignments for 30 individual students in a single semester. Generative AI changes that calculation in three specific ways.

Assignment Deconstruction

Most assignments bundle learning objectives together with a specific task and context. A research paper on a historical figure is, underneath its instructions, asking students to evaluate sources, construct arguments, synthesize evidence, and cite accurately. Those objectives can be reached through many pathways. A large language model can separate what an assignment is asking from how it is asking it, making the underlying objectives explicit. That deconstruction step is foundational. Once the objectives are clear and confirmed by the instructor, personalization can happen without losing the assignment's purpose.

Student Context Integration

Once an instructor has a set of student intake responses, a language model can map individual student interests, goals, and contexts onto the extracted learning objectives. A student in a sociology course who works in healthcare gets the same analytical framework applied to that domain. A student interested in immigration policy gets the same framework in a different context. The cognitive demand is equivalent. The entry point is the student's own.

The intake questions are designed to reach into a student's funds of knowledge and lived aspirations without requiring disclosure of deeply personal material. They ask what students value and are working toward, where they already have knowledge or confidence relative to the course content, and what communities and problems feel personally significant. That is enough to build a genuine entry point.

Rubric Consistency

A language model can generate a rubric that assesses learning objectives independently of the specific content pathway. Rather than evaluating whether a student wrote about a particular case, the rubric evaluates whether they demonstrated the analytical operation the assignment was designed to test, in whatever domain they applied it. An instructor can grade 30 different assignments against the same criteria. The evaluation is consistent. The paths to it are not.

VI. The Proposed Workflow

The workflow that follows is designed around a core principle: the instructor drives every decision. AI supports analysis and generation at scale. The instructor verifies before anything moves forward. The framework does not ask instructors to abandon their assignments or their expertise. It asks them to let those things do more work by routing them through what students already know and care about.

Each phase below is described in prose and summarized in the table that follows. The SDT and CRP grounding column in the table names exactly what need or mechanism each step is designed to serve.

Phase 1: Instructor Submits.

The instructor uploads their existing assignment to the tool. No redesign happens yet. The tool reads the assignment as written. This step is intentionally minimal because the instructor's original work is the foundation, not a problem to be fixed.

Phase 2: Instructor Verifies.

The AI deconstructs the assignment, separating the underlying learning objectives, cognitive demands, and skills being tested from the surface-level task instructions. It also surfaces the implicit cultural assumptions embedded in the assignment context. The instructor reviews this analysis and confirms or corrects it before anything else proceeds. This is the most important decision point in the workflow. The instructor's pedagogical intent is locked in here and protected through every step that follows.

Phase 3: Instructor Distributes.

Using the verified learning objectives, the AI generates a set of student intake questions. These questions are calibrated to surface what students value and are working toward, where they already have knowledge or confidence, and what communities and contexts feel personally significant. The instructor reviews and edits the questions before distributing them through whatever intake mechanism fits the course, a learning management system, Google Forms, or Microsoft Forms. The instructor decides what gets asked and how.

Phase 4: Students Respond.

Students complete the intake on their own time. The tool includes a brief note to students clarifying that no personally identifiable information is needed and that responses will be used only to shape their assignment, not stored in any official record. Students describe their interests, goals, and contexts in general terms. They are not required to disclose anything they are not comfortable sharing.

Phase 5: Instructor Reviews and Releases.

The instructor uploads the batch student responses. The AI maps each student's context onto the verified learning objectives and generates one individualized version of the assignment per student. It also generates a single shared rubric that evaluates the learning objectives in domain-neutral terms. The instructor reviews all individualized assignments before students see them. Any assignment that does not meet the instructor's standard goes back for revision. Nothing reaches students without instructor approval.

Phase 6: Students Complete.

Each student receives their individualized assignment. The intellectual demand is the same across all versions. The context is their own. The goal is for the work to connect to something the student already carries, so that engagement is organized from within rather than enforced from without.

Phase 7: Instructor Grades.

The instructor applies the shared rubric to all student submissions. Because the rubric assesses the cognitive operation and not the specific domain, it works consistently across assignments that look different on the surface. Every student is evaluated on equivalent criteria. The paths to meeting those criteria are different. The standard is not.

Workflow Summary

Phase 1 *Instructor Submits*	**What happens**	**SDT / CRP grounding**
	Instructor uploads existing assignment. No redesign yet. Original work is the foundation.	Instructor authority preserved from the start. The assignment's purpose drives everything that follows.

Phase 2 *Instructor Verifies*	**What happens**	**SDT / CRP grounding**
	AI deconstructs learning objectives and surfaces cultural assumptions. Instructor reviews and confirms before proceeding.	CRP: Reveals whose cultural context the assignment was designed around. SDT: Locks in the instructor's pedagogical intent before personalization begins.

Phase 3	**What happens**	**SDT / CRP grounding**

Instructor Distributes	AI generates student intake questions. Instructor edits and approves before distributing.	SDT: Questions target autonomy (values, goals), competence (existing knowledge), and relatedness (communities, contexts). Instructor controls what is asked.

Phase 4 *Students Respond*	**What happens**	**SDT / CRP grounding**
	Students complete intake with privacy protections. No PII required. General context only.	SDT: Supports autonomy by letting students choose what to share. CRP: Reaches into funds of knowledge without requiring personal disclosure.

Phase 5 *Instructor Reviews and Releases*	**What happens**	**SDT / CRP grounding**
	AI generates individualized assignments and shared rubric. Instructor approves before students see anything.	SDT: Each assignment satisfies autonomy, competence, and relatedness for that student. CRP: Entry point drawn from student's own cultural schema. Instructor retains full authority.

Phase 6 *Students Complete*	**What happens** Each student receives their personalized version. Same intellectual demand. Different entry point.	**SDT / CRP grounding** SDT + CRP: Work connects to something the student already carries. Conditions for negentropy are in place.

Phase 7 *Instructor Grades*	**What happens** Instructor applies shared rubric to all submissions. All students evaluated on equivalent criteria.	**SDT / CRP grounding** Rubric assesses the cognitive operation, not the content pathway. Consistent evaluation across individualized assignments.

VII. The Equity Argument

Students who arrive at higher education with cultural capital aligned to the dominant contexts of academic assignments already have a version of relatedness built in. An assignment drawing on Western canonical literature, corporate financial structures, or standard academic English conventions connects to the lived experience of students whose prior schooling prepared them for exactly those contexts. This is not accidental. It reflects whose knowledge has historically been treated as the default.

Students without that alignment are not less capable. They are less scaffolded. I work with these students in open-enrollment postsecondary settings. A student who is articulate, analytically sharp, and clearly capable may disengage from an assignment not because the intellectual work is beyond them, but because the context the assignment assumes has nothing to do with their life. The failure is in the design.

Walton and Cohen's (2011) research on belonging uncertainty is instructive here. Their work documents that students of color and first-generation college students disproportionately carry a chronic concern about whether they belong in academic settings, and that this uncertainty undermines performance independent of ability. An assignment that treats a student's context as a legitimate intellectual entry point does not simply engage them. It sends a signal that their knowledge counts. That signal is an equity intervention at the level of instructional design.

AI-supported personalization raises the floor of accessibility by ensuring every student has an entry point drawn from their own experience and schema. The students who benefit most are the students whose contexts have been least represented in conventional academic assignments.

There is a secondary effect worth naming. A personalized assignment that routes genuine intellectual work through a student's own frame of reference may be harder to delegate to AI, because it requires the student to engage with their own context. The approach addresses the motivation condition that produces AI workarounds rather than chasing the workarounds themselves.

VIII. What This Is Not

A few misreadings are worth addressing directly. This is not a criticism of how assignments have been designed. The assignments instructors have built reflect real pedagogical intentions, and those intentions are preserved through this process. What changes is the surface-level context in which the intellectual demand is situated, not the demand itself.

This is not a system for generating easier assignments. A student applying sociological theory to sports media is doing the same intellectual work as a student applying it to healthcare.

This is not a replacement for instructor judgment. The instructor confirms the AI's deconstruction of their assignment before any personalization proceeds. The instructor retains full authority over how student work is evaluated. The AI handles pattern-matching at scale. The instructor handles meaning.

IX. Conclusion

Four in 5 university students now use generative AI. Formal education is lagging behind. Those two facts describe a gap that detection tools will not close. The gap is motivational and cognitive, and it falls hardest on the students whose educational histories have already left them with the least margin.

SDT tells us what students need psychologically to engage with genuine investment. Hammond tells us why certain students have been systematically denied those conditions and what that costs neurologically. Walton and Cohen tell us that belonging uncertainty is not a side effect of academic environments but a structural condition that shapes whether students can access their own capacity. Generative AI provides the mechanism that makes a response to all three frameworks possible at the scale higher education requires.

The institutions that respond to the AI moment by asking only how to detect and restrict will spend the next decade chasing a moving target. What makes academic work worth doing for the students?

References

Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry, 11*(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01

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

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

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.

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

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., & Cohen, G. L. (2011). A brief social-belonging intervention improves academic and health outcomes of minority students. *Science, 331*(6023), 1447–1451. https://doi.org/10.1126/science.1198364