Formative AI Pedagogy: Designing for Depth in an Age of Efficiency

L. Callid Keefe-Perry*

Abstract

This article proposes a formative AI pedagogy for theological education, rejecting the use of technology for mere efficiency. It introduces the ANCHOR framework to ensure AI tools extend student capacity, scaffolding deep learning and vocational formation rather than replacing human agency and meaning-making.

Keywords: Formative AI pedagogy, ANCHOR framework, Capacity extension, Self-determination theory

It began with a modest experiment. Last year, during the opening session of our field education course, I introduced a prototype artificial intelligence (AI) tool I had clumsily named the Theological Lens Finder.[1] It was an early iteration of a tool I would later refine, and I offered it to the class not as a requirement but as an invitation—a digital conversation partner to help students just starting theological education, broadening the resources they had to draw on to do theological reflection. They could type or talk to the tool, and it would suggest possible lenses that might provide them with some insight. I demonstrated how to use it to deepen their inquiry, encouraged them to experiment, and noted it was available to use when they were struggling to come up with a fitting theological lens or perspective to use in their theological reflection assignments.

A few weeks later, as I was grading the second batch of theological reflections, I paused over one submission. The reflection itself was nuanced, vulnerable, and theologically rigorous. But what stopped me was a small block of text on the middle of the final page, a post-script typed out after the end of the formal paper: "P.S. How can this ChatGPT Theological Lens Finder be so kind and validating?!? Its responses spoke to the depths of my heart."

I thought the student was being humorous and mocking, but the following week I saw the student and mentioned the note. We stood in the hallway after class, and they confessed to a strange internal conflict. They knew that the "kindness" they experienced

* L. Callid Keefe-Perry is Assistant Professor of Public Theology and Director of Contextual Education for Boston College's Clough School for Theology and Ministry, Boston, Massachusetts. Email: callid.kp@bc.edu.

Reflective Practice: Formation and Supervision in Ministry, volume 46 (2026)
ISSN 2325-2855 (online)

wasn't real per se. Yet the experience of being heard felt undeniable. They told me they had begun using the tool for personal reflection outside the bounds of the assignment. I explained that the tool's "kindness" was architectural, patterned on my own pastoral voice and affirming feedback style. I told the student, "Think of it not as a replacement for me but as an extension of my presence." It was in that hallway conversation that the true shape of this pedagogy clicked into place.

I realized the tool extended my availability rather than replacing my labor; it was about presence, not efficiency. This conversation shifted my thinking. I had already been working at using AI, but after that day it was clear that the task before us was not merely technical but deeply pedagogical. I was no longer just asking how to make the tool work; I was asking how to tend to the student's formation in a digital space.

INTRODUCTION

To understand how such a "hallway moment" is possible with a machine, we must distinguish between the default use of AI and the approach I explore in this article. Used "as is," a standard large language model (LLM) mostly functions like a hyper-efficient search engine, which can seduce overworked students. My approach, however, customizes AI specifically for the kinds of learning I want students to experience rather than plopping "out-of-the-box" AI into the classroom.

Questions about formation in digitally saturated environments are not new topics for this journal. Rolf Nysse's work on "digital natives" and the formation of emerging religious leaders and Stephanie Paulsell's reflection on technology and ministry both named how shifts in media reshape the habits and expectations of pastoral formation.[2] This essay picks up that thread in relation to AI, asking what kinds of AI use can deepen, rather than thin out, the work of theological reflection and supervision.

I am intentionally omitting the technical mechanics of building such tools here to focus on the pedagogical framework. However, it is important for readers to know that moving beyond generic AI does not require coding expertise. Customizing an AI for the classroom is now achievable by typing clear, natural-language instructions, much like training a teaching assistant. You define the learning goal, provide the reading list, specify the tone, and set boundaries for what the tool should not do (e.g., "Do not write the student's paper"). The barrier to entry is no longer technical know-how but pedagogical clarity.

While the arrival of AI rightfully makes us cautious about tools that might shortcut thinking, when designed with this kind of intentionality, AI need not compete with our pedagogical goals. The skill required is not computational; it is rhetorical.[3] If any reader is interested in exploring and experimenting with using AI this way, I encourage them to reach out; I'd love to support other educators in developing their skills in this regard. In the interim, however, the focus for this article is on the *orientation* we need as educators

considering the use of AI, not the specific technique of that implementation. A formative AI pedagogy uses technology to support human growth, not replace it.

PRINCIPLES FOR FORMATIVE AI PEDAGOGY

To define what is at stake, we might distinguish between the mere completion of academic tasks and what the educational philosopher Chris Higgins calls "formative education." Higgins defines this not as a transactional exchange but as "a broader existential task, the quest to understand, cultivate, and enact ourselves in lives worth living."[4] A formative AI pedagogy, therefore, is one that orients the student toward this existential task. It rejects the "transactional" model where education is simply the efficient production of outputs. Instead, it asks a fundamental design question: "How can the interaction with this technology help the student become the kind of person capable of doing this work?" What, then, does a formative approach to AI actually look like in practice? It requires a deliberate pedagogical shift; we must resist the technology's bias toward speed and volume, pivoting instead toward depth and interiority.

A formative approach does not use AI merely for information processing or efficiency; rather, it scripts the technology to support meaning-making, autonomy, and deep learning. As I have been thinking through this work, I've developed what I call the ANCHOR Framework for formative AI pedagogy, intended to help do just that. This approach evaluates AI integration through six distinct principles: **A**gency, **N**eighbor-Connectedness, **C**ompetence, **H**ermeneutics, **O**utward Orientation, and **R**eflection. The framework centers on a fundamental distinction between capacity replacement and capacity extension.

Replacement occurs when AI performs the core task (e.g., generating reflections), atrophying the student's cognitive muscles. Extension uses AI to scaffold abilities, pushing students deeper into the productive struggle of learning. The ANCHOR principles serve as a guide to ensure our tools always lean toward extension, preserving the student as the primary architect of meaning. The first three principles (ANC) draw from Deci and Ryan's self-determination theory (SDT),[5] while the final three (HOR) expand into vocational formation. This dual grounding evaluates whether a tool or exercise fosters growth or merely processes inputs.

The first three principles of this framework distinguish between controlled motivation, which is driven by external rewards or punishments, and autonomous motivation, which arises from a sense of volition and personal value. Deci and Ryan argue that for individuals to experience genuine growth, integration, and well-being, three basic psychological needs must be met: autonomy (the need to feel that one is the origin of one's own actions), competence (the need to feel effective in interacting with the environment), and relatedness (the need to feel connected to and cared for by others).[6]

For theological educators, SDT offers a compelling empirical basis for what we might otherwise call "formation." While traditional educational models often leverage

extrinsic motivators such as grades, degree requirements, or institutional compliance to drive student performance, SDT suggests that deep, enduring learning requires the internalization of these values. When the learning environment supports autonomy, competence, and relatedness, students move from mere compliance to intrinsic motivation. They engage in the work not because they must but because the work itself has become meaningful to their identity and vocational trajectory.

While external pressure acts like a weight that temporarily bends a branch, formative pedagogy nurtures the learner's interiority so they grow in a new direction from the inside out. This ensures change remains even after external structures, like deadlines, fall away. In our current technological context, SDT provides a necessary critical lens. AI engines are designed for efficiency and output, attributes that naturally align with extrinsic performance goals rather than intrinsic growth. Without a pedagogical framework that prioritizes the internal conditions of the learner, AI tools risk functioning merely as "weights" on the branch, producing polished outputs without requiring the internal struggle that constitutes formation. By anchoring our AI pedagogy in the specific needs for autonomy, competence, and relatedness, we ensure that these powerful tools serve to scaffold the student's internal development rather than replacing the human effort that makes learning durable.

If the first three ANCHOR principles address the psychological engine, the final three provide the theological steering. These shift the focus from psychological well-being to communal responsibility, yoking the student to the service of their neighbor. They pull the student out of their own head and situate them within a wider web of accountability. Hermeneutics tethers the student to a historical community and tradition; Outward Orientation directs their agency toward the present needs of the neighbor and the pursuit of justice; and Reflection ensures they are honestly assessing their internal formation in light of these external demands.

Agency

In SDT, Agency maps to the need for autonomy—the experience of volition where actions are self-endorsed rather than externally pressured. Fostering agency shifts students from passive recipients who are complying with authority to active authors, increasing persistence and their capacity to make meaning.

When a curriculum is overly prescriptive, dictating every step and interpretation, students may perform the tasks, but they often fail to internalize the skills. AI can make this much worse. Good pedagogy, therefore, introduces "autonomy-supportive" structures: offering meaningful choices in assignments, encouraging self-initiation, and providing a rationale for constraints. When students feel ownership over the process, they are more likely to engage in the kind of deep, reflective struggle that leads to lasting growth.

In the context of AI, prioritizing agency becomes urgently important. The default design of many generative AI tools encourages passivity; the user asks a question, and the machine delivers a complete, polished answer, effectively removing the student from the decision-making loop. Without intentional pedagogical intervention, AI can easily become a "crutch" that bypasses human agency entirely. A formative AI pedagogy counters this by designing interactions that require the student to remain the primary actor. By scripting the AI to wait for the student's lead, we increase the likelihood that the technology amplifies human agency rather than erodes it.

Neighbor-Connectedness

In SDT, this principle maps to relatedness, the deep need for reciprocal care and community belonging that drives the internalization of group values. Framing this as Neighbor-Connectedness expands this psychological concept into an ethical mandate. It emphasizes that true formation requires moving beyond mere interpersonal attachment to recognize our well-being as bound up with the well-being of the "neighbor."

We know that learning is fundamentally a social endeavor. Good pedagogy constructs a "community of inquiry" where students feel safe enough to take risks and challenged enough to grow. When students see their learning as a contribution to a community rather than a private transaction for a grade, their motivation shifts. They engage more deeply because they understand that their formation matters not just to themselves but also to the peers and future communities they will serve.

In the age of AI, this emphasis on connection is critical because the technology carries an inherent risk of isolation. A student can easily spend hours in a loop with a chatbot, perfecting a paper or refining an idea, while entirely bypassing the friction and richness of human relationship. A formative AI pedagogy must actively resist this siloing effect. I do not want us to use AI to replace human conversation; it should instead prepare us for conversation. We can design AI interactions that function as on-ramps to community. For instance, we can develop a tool to help a quiet student organize their thoughts so they feel confident enough to speak in seminar, or use AI to simulate diverse perspectives that the student must then engage with in real life. The goal is always to use the tool to return the student to the neighbor, better prepared for the work of shared life and ethical responsibility.

Competence

In SDT, competence is the need to feel effective and experience true mastery through "optimal challenges." This aligns with Vygotsky's "zone of proximal development," where scaffolding satisfies the need for competence and fuels the motivation to build demonstrable skill.[7] The primary way educators can support competence is through scaffolding.

Good teaching does not simply assign a high-stakes task and hope for the best; it breaks complex skills down into manageable steps, providing models, feedback loops, and opportunities for low-stakes practice. This approach honors the value of "productive struggle."[8] We know that learning requires a degree of friction; the effort of retrieving an idea or refining a sentence is actually what encodes the knowledge. If a task is too easy, the student learns nothing; if it is too hard, they retreat into helplessness. Pedagogy is the art of keeping the student in that productive middle ground where growth happens.

In the context of AI, the principle of competence faces its greatest threat and its greatest opportunity. The threat is obvious: generative AI can instantly remove the struggle entirely. By producing a finished essay or a polished reflection in seconds, the tool bypasses the very friction required for learning, leaving the student with a product but no increased capacity. This is a "competence trap" where the machine is capable, but the user is not. A formative AI pedagogy must rigorously resist this. Instead, we treat AI as the ultimate scaffolding tool. Because it is patient and infinitely customizable, AI can be designed to operate within a student's zone of proximal development, offering a hint rather than an answer, critiquing a paragraph rather than rewriting it, or role-playing a difficult pastoral conversation. When used this way, AI does not replace the struggle; it makes the struggle productive, helping the student climb the ladder of readiness one rung at a time.

Hermeneutics

The move from information to hermeneutics is the threshold of deep learning. It is the shift from asking "What does the text say?" to asking "What does this mean for us, here and now?" Education that stops at information transfer is mostly transactional. Formative pedagogy, by contrast, prioritizes meaning-making. It challenges students not just to aggregate data but to synthesize it, contextualize it, and bring it into dialogue with their own lives and traditions. In theological education specifically, this is the heart of the discipline; we are training students not to be encyclopedias but to be wise interpreters who can discern significance amidst complexity.

In the context of AI, this principle is a necessary corrective to the technology's inherent bias toward surface-level summary. Because unmodified LLMs generate "smooth" answers and often flatten nuance, they risk encouraging a banking model "pedagogy of retrieval." A formative approach resists this by positioning AI as the starting point, rather than the conclusion, of interpretation. We might use AI to summarize a complex topic, but only so the student can spend their energy critiquing that summary or applying a specific theological lens to it. By assigning the machine the work of processing information, we free the student to take up the higher-order work of discerning meaning.

Practically, this can be implemented through a "hermeneutical friction" exercise. Here, students prompt the AI to interpret a case study through conflicting frameworks,

such as "prosperity gospel" versus "liberation theology." Forcing multiple plausible meanings shifts the student from a passive recipient to a primary interpreter who must decide which lens is most faithful to their tradition. This increases the likelihood the AI remains a provocateur for deeper thought rather than a surrogate for it.

Outward Orientation

One meaningful test of formation is not what the student retains but what they contribute. Education that turns the student inward, focusing solely on individual achievement or personal accumulation of knowledge, fails the fundamental vocational mandate. A formative pedagogy insists that the classroom is not a sealed container; it is a rehearsal space for public life. Consequently, every act of instructional design should be crafted in view of larger goals oriented not just to the self but also to the common good. The goal is to shape leaders who see their intellectual growth as a resource for their communities, ensuring mastery is always understood as preparation for service.

While Neighbor-Connectedness helps the student feel seen within a relational web, Outward Orientation ensures that this web is not a closed circuit but a conduit for public service. In the context of AI, this outward turn is a countercultural move. AI tools are overwhelmingly marketed as engines of personal productivity and individual advantage. They are sold as "co-pilots" designed to make our work faster and our life easier. If left unexamined, this reinforces a solipsistic loop where technology serves only the user. A formative AI pedagogy breaks this loop by explicitly directing the tool's output toward the needs of the neighbor. We might ask students to use AI not to write a better essay for a grade but to help draft an advocacy letter for a local nonprofit, simplify complex medical information for a parishioner, or design a resource for a youth group. By insisting that the technology be used to solve problems beyond the student's own academic survival, we tether the power of AI to the pursuit of the common good.

This can be operationalized through a "preparation for presence" activity. Rather than using AI to polish an internal reflection, students are tasked with using the tool to simulate a dialogue with a specific, challenging stakeholder they are scheduled to meet with, perhaps a skeptical city council member, a grieving parishioner from a different cultural background than their own, or a community organizer with conflicting priorities. By framing the AI as a sparring partner for public ministry, the technology ceases to be a tool for individual efficiency and becomes a site of vocational rehearsal, ensuring that when the student finally closes their screen, they are better prepared to offer their full, unmediated presence to the people they serve.

Reflection

Reflection is a necessary counterpart to hermeneutics. While hermeneutics focuses on interpreting the text or the world, asking, "What does this mean?," reflection turns the interpretive lens inward, asking, "What does this mean for *me*?" and "How is this

changing the way I think?" A student can brilliantly interpret ideas while remaining opaque to themselves. Formative pedagogy resists this by prioritizing metacognition, the ability to think about one's own thinking. Requiring students to examine their own assumptions ensures education cultivates an examined life rather than just the accumulation of data.

In the context of AI, reflection is the primary defense against the technology's bias toward speed. While AI accelerates output, formation requires the opposite: slowing down and dwelling with uncertainty. We might require students to submit a "process audit" alongside their AI-assisted work, detailing not just what the AI produced but also how they interacted with it, where they pushed back, and what they learned about their own limitations in the exchange. By prioritizing this metacognitive layer, we ensure that the efficiency of the machine does not hollow out the interiority of the student.

You could lean into this principle through something like a "prompting log." Instead of merely turning in a final product, students are required to submit a short narrative of the productive struggle they encountered during the session. They might be asked to identify the specific moment when the AI's response felt too smooth or misaligned with their own pastoral intuition and then describe how they redirected the tool to better serve their inquiry. By making the process of interaction a graded component of the assignment, the AI session is transformed from a transactional errand into a site of self-observation. This ensures that the machine's efficiency serves as a whetstone for the student's own critical interiority, preventing the work of metabolizing experience from being outsourced to an algorithm.

The ANCHOR Goal

Taken together, the ANCHOR framework offers more than a theoretical critique; it provides a design specification for the responsible integration of technology. If we leave AI tools at their default settings, don't design with AI in mind, and don't talk to our students about how to engage with it well, usage of AI will often drift toward efficiency, speed, and the displacement of human effort. However, by intentionally constraining our tools with these six principles, we can bend the technology back toward the goals of formation. To help navigate this, I propose a simple diagnostic distinction between "capacity extension" and "capacity replacement."

Capacity extension uses AI to scaffold a student's reach (organizing thoughts or critiquing drafts) while the student remains the architect of meaning. Conversely, capacity replacement substitutes AI for the student's own effort, resulting in a finished product but diminished internal growth. Here, the AI generates the insight, structures the argument, or provides the theological application. While the student may leave the interaction with a finished product, they do so with a diminished capacity for agency, having successfully integrated their task into a system that values output over interior growth.

To use this framework practically, educators must evaluate the entire pedagogical ecology, not just the tool itself. Before integrating an AI exercise, an instructor should ask: Does this specific use case require the student to engage in the productive struggle of learning, or does it bypass that necessary friction? If an application leans toward replacement, the ANCHOR principles suggest intervention. This might mean adjusting the AI's underlying instructions so it only suggests alternative lenses, but it equally requires explicit communication with students about the tool's intended boundaries and the value of their own meaning-making. Furthermore, it necessitates carefully placing the exercise within the semester's arc, ensuring it serves as a temporary scaffold rather than a permanent crutch. In this way, the framework operates as a comprehensive compass, shaping not just the technology but the entire learning environment to support cognitive and spiritual development.

In my classrooms, I try to be explicit about this distinction. If the AI is doing the "heavy lifting" of meaning-making, it is acting as an instrument of conformity; if it is helping the student lift the weight, it is serving their own growth and engagement with the complexities of the world. Translating this ideal into everyday classroom practice, however, requires educators to anticipate and address a series of ethical and pedagogical challenges.

NAVIGATING PITFALLS AND ETHICS

Adopting a formative AI pedagogy requires more than enthusiasm; it requires a sober assessment of the risks. In theological education, the line between using a tool to enhance capacity and using it to erode agency is perilously thin. If we are not careful, the very efficiencies that make AI attractive can lead to conformity, shaping students who are technically proficient at producing text but spiritually and intellectually hollowed out. Navigating this landscape requires us to attend to at least three potential ethical hazards.

The first is the hazard of resource stewardship. Digital intelligence carries heavy material costs in water and energy, an ecological burden demanding we move beyond mindless usage. However, we must also place this impact in context. While AI data centers are resource intensive, their global footprint is a fraction of sectors like agriculture or transport. For instance, based on current estimates, a 1% improvement in global agricultural water efficiency could easily save more than 10 km^3 of water annually, whereas global AI data-center water use is projected to reach only around 1 km^3 per year by the late 2020s.[9] The ethical task, therefore, is not to retreat into a paralyzing "prompt guilt" but to practice stewardship. If we are to draw on these resources, we must ensure the usage is directed toward vocational formation that justifies its ecological cost through genuine human growth.

The second hazard is that of mistaking the product for the process. Bypassing the messy, frustrating work of finding the right words for an experience isn't just an academic shortcut; it short-circuits the student's own formation. The student must instead be

challenged to critically evaluate the impact of the AI-assisted theological lens. The task isn't just for the student to identify an interpretive lens that works in general but to sit with a consideration of what that lens does to the student, to allow the theological insight to interrogate their own assumptions. The act of interpretation and theological reflection should be more of a prayerful practice of self-implication than one of abstraction and correspondence.

To counter this, educators must be clear in their communication. Custom tools can be explicitly scripted with a "Do not draft" directive, enforcing the distinction between assistance and substitution. We can insist that AI be used for retrieval and feedback, never for composition. The guidance is simple: the AI can carry the library, but the student must carry the pen. In my own course policy that I share with students, I frame this boundary not as a matter of policing but as an invitation to discernment:

> The core of your Theological Reflection papers and your concluding Essay on Ministry and Call is the work of personal reflection, a process of listening to your own life and bringing it into conversation with our course materials. You are expected to submit your own original, personally authored work for these assignments. Submitting text generated by AI as your own reflection is not permitted because it circumvents the primary goal of the assignment, which is your own spiritual and theological formation. Could you use AI to write a reflection that you didn't actually reflect on? Yes. Would I always be able to tell? No. However, the purpose of graduate theological education is to form reflective, discerning ministers. This course is built on the understanding that you are here to engage in the deeply personal work of vocational discernment. Entrusting the work of reflection to AI is an abdication of that sacred, co-creative responsibility. Use AI to help you do the work, don't replace your agency with the labor of AI.
>
> I am trusting you to engage the assignments with integrity, not because I will be policing your every word but because the work of theological reflection is essential to your own vocational journey. Choosing to do the reflective work yourself is an act of faithfulness to your own formation. It is how you develop the "discernment muscles" necessary for a life of ministry. Ultimately, you will get as much out of this class as you put into it, and my prayer is that you will choose the path that leads to your own deepest growth and flourishing.

Does this guarantee that an enterprising and unscrupulous student won't figure out a way to circumvent these protocols? No. But the goal of formative education is not to build a surveillance state; it is to build a conscience. If we reduce pedagogy to a game of cat-and-mouse, we reinforce the very "controlled motivation" that undermines deep learning. By framing the decision to write as a vocational imperative rather than an arbitrary rule, we move the locus of control from the professor's oversight to the student's interiority, leaning into the critical hope that the desire to be formed is a more durable constraint than any technological barrier.

The third hazard is artificial intimacy. To mitigate the risk of artificial intimacy, we must demystify the "magic" through radical transparency about the tool's design. A

formative AI pedagogy requires that we pull back the curtain on why the AI sounds the way it does. We need not always have students read the raw system prompt, but we must be explicit about the architectural constraints that govern the interaction. For instance, in the hallway conversation mentioned earlier, I clarified to the student that the "kindness" they experienced was not sentient but structural: the model had been specifically patterned on my own language for giving feedback and instructed to simulate best practices in compassionate spiritual accompaniment. By revealing these pedagogical inputs, the interaction shifts from an interpersonal encounter to an appreciation of design. The student realizes they were not "seen" by a ghost in the machine but were supported by a deliberate extension of the teacher's pedagogical and pastoral intent. This not only clarifies the nature of the tool but also, ideally, helps them to think about how they might design similar tools for their own context(s).

CONCLUSION

The integration of AI into education is not a question of if but how. Currently, the default affordances of these tools often stand in tension with the slow, interior work of deep learning. The ANCHOR framework is offered here not as a prescriptive or exhaustive solution but as a heuristic, a way to bend technology toward formation rather than accepting its default settings. Beyond technical considerations like retrieval-augmented generation, the core questions are fundamentally pedagogical. They require a deliberate decision to value the productive struggle of learning over the efficiency of the output. As this paper has argued, the distinction between a tool that hollows out agency and one that extends it does not lie solely in the code but is also found in the educational commitments that surround it. It is the choice to design interactions that use technology to *scaffold* the effort necessary for growth rather than bypassing complexity.

In the opening vignette, the student in the hallway experienced a pedagogical opening where pastoral care was mediated through a digital proxy. However, we must remember this tool is a functional simulacrum of kindness, designed to mirror, not replace, my educational intent. Mistaking artificial intimacy for a final destination fails our students. Simulated kindness must only serve as a bridge, scaffolding internal stability so students can eventually offer their full presence to the world. If the machine becomes a destination rather than a bridge, we have succumbed to the logic of efficiency. The aim of formative AI pedagogy is to return the student to their community, not with a generated answer but with a widened imagination and a deepened capacity for listening. We should ensure that the student remains the architect of their own meaning, using the tool to scaffold the productive struggle that leads to genuine, integrated growth.

We must not be afraid of AI, nor should we surrender to it. Instead, we must steward it. By anchoring AI in the specific human needs for agency, connection, and interiority, we can bend the arc of technology toward formation. In doing so, we ensure that our students emerge from the classroom not merely as efficient operators of a system

but as wise, empathetic, and fully formed practitioners capable of guiding others through the productive struggles of their own formation.

NOTES

[1] The current version of the tool is free and is available at http://wrentool.app. WREN (Wisdom Reflection Explorer & Navigator) is an AI tool that utilizes retrieval-augmented generation (RAG) to provide students with a formative, rather than transactional, AI interface. It is designed to act as a conversational partner in the theological reflection process.

[2] Rolf Nysse, "Learning from 'Digital Natives': Forming a New Generation of Religious Leaders," *Reflective Practice: Formation and Supervision in Ministry* 31 (2011): 11–19; Stephanie Paulsell, "Technology and Ministry," *Reflective Practice: Formation and Supervision in Ministry* 31 (2011): 20–34.

[3] The rhetorical skill I am referring to is found not in complex coding but in the precision of the pedagogical directive. It is the difference between a student asking an AI tool to "summarize this text" and an instructor providing an AI tool that has been told to help a student "interrogate this text for hidden biases while refusing to provide a summary." By engineering these intentional constraints, the educator transforms a tool of efficiency into a scaffold for productive struggle.

[4] Chris Higgins, *Undeclared: A Philosophy of Formative Higher Education* (MIT Press, 2024), ix.

[5] Edward L. Deci and Richard M. Ryan, *Intrinsic Motivation and Self-Determination in Human Behavior* (Plenum, 1985).

[6] More about SDT and its application to education can be found here: "Applying Self-Determination Theory to Education," Center for Self-Determination Theory, https://selfdeterminationtheory.org/topics/application-education/.

[7] Richard M. Ryan and Edward L. Deci, "Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being," *American Psychologist* 55, no. 1 (2001): 68–78.

[8] Jamaal Young, Danielle Bevan, and Miriam Sanders, "How Productive Is the Productive Struggle? Lessons Learned from a Scoping Review," *International Journal of Education in Mathematics, Science, and Technology* 12, no. 2 (2024): 470–95. https://doi.org/10.46328/ijemst.3364.

[9] Pengfei Li, Shaolei Ren, Dandan Li, and Zhi Zhou, "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models," last revised March 26, 2025, https://arxiv.org/abs/2304.03271; "AI Data Centers to Drive 11-Fold Rise in Water Consumption by 2028: Morgan Stanley," *The Economic Times*, September 7, 2025, https://economictimes.indiatimes.com/tech/artificial-intelligence/ai-data-centers-to-drive-11-fold-rise-in-water-consumption-by-2028-morgan-stanley/articleshow/123755912.cms.