

REPORT OF THE AI TASK FORCE

Making AI Work for the Good of Our Students

Introduction

“Artificial intelligence,” or “AI,” has been an explicit goal of leading computer scientists since the middle of the 20th century, with philosophical antecedents reaching back at least to Descartes and arguably back to Aristotle’s systematic treatment of logic. Despite this long academic lineage, most non-specialists could serenely ignore AI until November 30, 2022, when a relatively little-known non-profit called OpenAI released the computerized conversational system known as ChatGPT. ChatGPT processes a user’s natural-language requests and generates responses that mimic the way a human interlocutor would respond—if the human interlocutor had somehow read nearly everything ever published.

It did not require much imagination to see that ChatGPT (and similar AI products that would quickly follow) might disrupt education, for better and for worse. Optimists imagined these systems providing students with personalized tutoring at scale—delivering lessons at each student’s pace, answering each student’s questions, assigning exercises aimed at each student’s level of understanding, and providing the correction and coaching that each student needs in order to gain in knowledge and understanding. Some even imagined these systems could replace teachers entirely. Others imagined that AI educational systems would make teachers far more knowledgeable and far more available to their students; that AI would help teachers quickly and almost effortlessly generate lesson material, create exercises and activities, and assign grades and other feedback on work submitted by each student. Ideally, this would free the teacher to focus more on the most human parts of the educational project, like motivation and mentoring.

Not everyone was optimistic, however, and it took even less imagination for pessimists to see what might go wrong. Some teachers, perhaps over-persuaded by the optimists’ hype, worried about their jobs. More soberly, everyone saw that these systems made plagiarism much easier for students to commit and harder for teachers to detect. It was easier to commit because the AI system could generate an entire essay that responded specifically to the teacher’s prompt, relieving a cheating student even of the bother of finding something relevant to plagiarize. And it was harder to detect because the text copied by the student was (in a way) original—it had been conjured into existence by the AI system, and therefore a suspicious teacher would not be able to identify the source by looking for it on the internet.

In addition, pessimists worried about whether these systems—and the tech companies that designed them—could be trusted to educate our children with the right substantive content. Google’s Gemini system, trying to one-up ChatGPT by producing images, amply illustrated this concern in February 2024, when user requests to generate a picture of Nazi soldiers produced a demographically

diverse group that included, for example, an Asian woman. A request for a picture of America's founding fathers similarly produced a picture that looked a bit more like the cast of *Hamilton*. Evidently, *someone* had programmed some substantive side constraints into the system that would generate false answers, instead of the objective and easily discoverable truth, if the truth were somehow undesirable (to *someone*). The Gemini failures that made the news were so egregious they could not fool any educated person, but this was supposed to be a product that educated the uneducated. And anyone could see that if the same or similar biases were also altering responses more subtly, these systems might not be educating as much as miseducating—depending on what the *someones* designing the systems wanted them to teach.

Veteran educators may be excused for thinking they have seen this all before—technologists making extravagant claims for the educational value of their technologies, from Smartboards to Schoology to Chromebooks or iPads for every student. But generative AI—which, for purposes of this Report, we can define as data-driven computational systems that can perform educational tasks normally requiring human judgment, including prediction, adaptation, feedback, generation, or recommendation—is already a much more powerful tool than any of the earlier New Things in ed-tech that turned out to be flashes in the pan. Moreover, generative AI is already reshaping other sectors of the economy, and the pace of technological advancement is quite fast and perhaps even accelerating. Probably no one in any occupation can afford to ignore it entirely; certainly teachers cannot ignore it at all. And indeed, teachers at The Heights are not ignoring it and have for several years been changing their pedagogy to accommodate the new reality (such as by not assigning take-home writing, and by not giving students the prompts for in-class blue-book essays the night before the exam).

In light of these developments, in December 2025, the Headmaster created the AI Task Force to study the use of AI to enhance student learning, and to make recommendations to the Administrative Council by the summer of 2026, “with an eye to implementing as many as possible and desirable in the 2026-27 academic year.” Accordingly, in the last six months, the Task Force has investigated the risks and opportunities presented by generative AI, and has tried to think both theoretically and practically about best (and worst) practices for AI use by teachers and students.

Executive Summary

- I. Generative AI does not change our mission, but it has already influenced our pedagogy and it will continue to do so, whether we like it or not. Fidelity to our mission will require our faculty to redouble our longstanding emphasis on the personal dimensions of our teaching and on the academic rigor we must require of our boys if we want them to grow in virtue and in competence in the liberal arts. This applies to every subject taught at The Heights.
- II. Generative AI has the potential to make our teachers even better at what they do, by amplifying each teacher's efforts to advance his own knowledge and to curate the rich intellectual tradition of the West for our students. Generative AI can even assist teachers with what is arguably the most intensely personal dimension of their teaching—the evaluation of our students' work. There are potential pitfalls, and we must beware of them, but we have no right to shrink from learning and using the best professional tools available to us, including tools based on generative AI. We hope this will become the subject of more professional development during faculty seminars.

- III. Student use of generative AI presents a very different cost-benefit profile, because of the substantial risk that some students will use the new tools in ways that reduce academic rigor and thus inhibit each student's intellectual development. We therefore recommend that generative AI products be presumptively off limits for all academic work at The Heights, with a hard and fast rule against using AI either to summarize assigned readings or generate text that the student presents as his own. There will be rare occasions on which an older and fairly advanced student's intellectual development may be promoted by allowing the student to use generative AI—research papers, for example—and we commit such exceptional cases to the sound discretion of the teacher, in consultation with his school head. In those rare cases, students and their parents must be apprised of the ways in which the overuse or misuse of AI will likely be not only dishonest or unfair but ultimately self-defeating. We hope this will become the subject of more formative discussion at one or more of the student assemblies that now occur one day per month.
- IV. Our survey of the current and potential uses of generative AI in education, at The Heights and elsewhere, has brought to our attention several closely related challenges on which we hope the Administrative Council will speak at the same time it addresses the use of generative AI. In particular,
- A. We are concerned that generative AI has already adversely affected writing instruction everywhere, including at The Heights, and this is especially regrettable for students in the middle of the class academically. We think this trend should be reversed, but doing so will probably require more substantial changes than the AI Task Force is competent to consider on its own. We recommend further discussion of these matters within the appropriate faculty groups.
 - B. One of the most effective ways to protect students from the harmful effects of taking the AI shortcut is to create assignments that need not and perhaps cannot be completed with the aid of a computer. Again, there will be exceptions, and this is a matter that must be committed to the teacher's sound professional judgment, in consultation with the head of his school.
- V. To our own surprise, we have come to see great potential, across numerous disciplines, in the supervised use of computer resources for remedial instruction and practice with foundational concepts and skills that students must acquire for proficiency at higher levels. We can take the multiplication table as our paradigm for this sort of obstacle to a student's further progress, because a student who cannot reliably multiply cannot reliably do algebra. But we believe similar gaps are present in other subjects, specifically including Latin, logic, grammar and composition, and perhaps even reading comprehension. Some of these computerized resources, like Math Academy, will be from third parties. But Heights teachers can now create more focused resources with the help of generative AI. We have referred to this internally as the Heights Learning Lab, and it would ideally be a staffed location where students in need of remedial help would be able to work large numbers of exercises targeted specifically at the deficiencies they are showing in class. To be clear, these resources are not themselves examples of generative AI, but they are resources that the vast majority of teachers would have been unable to create without generative AI. With generative AI, they are fairly easy to create, and they may provide a way for teachers to help struggling students more effectively despite the ordinary time constraints on both teacher and student.

- VI. Formally, we recommend that the student handbook be revised to include, in substance, the following (under “Academic Honesty”):

A student may commit an offense against academic honesty by:

1. cheating on a quiz or test;
2. copying another student’s homework;
3. plagiarizing, or presenting as one’s own the work of another author, including any unattributed or misattributed use of text generated by an AI platform or product; or
4. using unauthorized aids, including the use of an AI platform or product to summarize, excerpt, or digest any assigned material or any work the student represents himself as having read.

I. AI Does Not Change Our Mission

Our school’s mission statement proclaims the purpose of a Heights education: what we do, for whom, and for what ends. It states, “Our mission is to assist parents in the intellectual, moral, physical, and spiritual education of their sons, with dedicated teachers training boys rigorously in the liberal arts. This formation in virtue fosters respect for every person, a desire to serve God and others, and an optimistic attitude towards life’s challenges.” We aspire to be *rigorous* in our training because we know that our students cannot become the men God wants them to be unless we challenge them properly. We also know that if we were to relax our insistence on academic rigor, the boys who would suffer most would not be the self-motivated few at the top of the class; it would be the larger number of perfectly ordinary students in the middle of the class who need more help from us in order to realize their potential.

Our website affirms our “commitment to foster the complete *human* development of every boy” (emphasis added), through *personal* instruction in virtue and in the liberal arts. We characteristically try to advance this goal by focusing on the agency of our students. We seek to make of them what Walker Percy described as “sovereign knowers,” a phrase [Michael Moynihan unpacks](#) to mean “a person who has governance of how he stands before reality, knowing in a way that takes ownership of the knowing, as one who has full possession of his reason, exercising it in the present moment. A sovereign knower is engaged and self-directed, committed to seeking to know truth for himself.” Our success in this endeavor depends not just on the soundness of our curriculum but on the teacher’s ability to teach [“in such a way that the student becomes engaged in learning, that he falls in love with learning and takes ownership of what he learns.”](#) Importantly for present purposes, one of the surest ways that teachers can frustrate this goal is by spoon-feeding students with information for them to learn superficially, and thus depriving them of the beneficial struggle that authentic learning always requires.

Given this mission, and this strategy for fulfilling it, there is no question of AI actually replacing teachers at The Heights. In-person instruction has always been the cornerstone of our pedagogical approach, and no one who remembers the year of COVID-19 accommodations could possibly want to change that. We rely on teachers not only to know the material they teach and to know how to teach it, but to know the students and know how to motivate them to achieve excellence—and excellence, not adequacy, is our goal.

Similarly, there is no question of permitting students at The Heights to use AI as a shortcut that bypasses the human element on their side of the student-teacher relationship. In some sense, the goal of the teacher is always to make learning easier for the student. The paradox, however, is that one of the most important ways we make learning easier for them is by assigning them work that is hard. If the work gets too easy, we inhibit their intellectual development. To put this another way: We value their hard work not for its own sake but for the sake of the transformation that only hard work can bring about. Accordingly, just as we emphasize the human dimension of the teacher's work, we also emphasize the human dimension of the student's work. To this point, [a recent study by Barcaui](#) showed traditional learners outperform AI-assisted learners by 11 percentage points on a surprise knowledge retention test 45 days after studying the material. One proposed reason for this outcome was that the AI tool facilitated a form of cognitive offloading that eliminated the desirable difficulties needed for deep learning and knowledge retention. As a result, education frameworks are needed to ensure that AI serves as a supplement (not replacement) for the desirable difficulties that foster durable learning.

The Task Force thus understands our mission to imply two strong presuppositions about generative AI: first, that we should embrace generative AI to the extent it helps teachers become better versions of themselves and teach more effectively; and second, that we should avoid, prevent, or mitigate student engagement with generative AI to the extent it threatens to undermine or impede the development of our students' natural intellects. In other words, we take our mission to be fixed and non-negotiable, and the question before us is whether generative AI has any role to play in our pursuit of that mission.

II. AI Presents More Promise than Peril for Teachers

During the last six months of Task Force meetings, we have tried to educate ourselves about the ways generative AI might be used and is already being used at The Heights. We have also read many reports and reflections on the use of AI in schools and other educational settings. In case after case, we have returned to the truism that there are both opportunities and dangers involved. Over time, however, we have concluded that the promise and the perils of using generative AI net out quite differently for teachers on the one hand and students on the other. For teachers, we believe generative AI can greatly enhance each teacher's ability to carry out at least three essential functions of good pedagogy: the teacher's role as student of a discipline (*i.e.*, the teacher's own research); the teacher's role as curator of an intellectual tradition (*i.e.*, the teacher's role in designing courses and delivering content); and the teacher's role as judge of each student's intellectual achievement (*i.e.*, grading). We therefore urge teachers to incorporate generative AI into their professional toolkit not to make their work easier but to make their work better. Teachers should use AI for the same reason they should do anything as teachers: because it can benefit their students.

There are, however, three important caveats: First, teachers should not use AI in any way that communicates to the students a lack of human regard for their efforts; that would be bad for the students. Second, teachers should not use AI in any way that dulls their own keenness of intellect; that would be bad for the teachers. And finally, we are writing with veteran teachers in mind, not rookies; a relatively new teacher who wants advice on how to become a better teacher should trust the same human instruction we provide for our students, relying primarily on mentors, not machines.

With these caveats, we survey the upsides and downsides of generative AI use by teachers.

A. Teachers as Students of Intellectual Disciplines

One of the great joys of teaching is the opportunity to remain a lifelong learner. The old saw about learning a subject best by teaching it to others is one we can all verify from our own experience. But that requires research, and the old-fashioned Google search (!) has now been almost completely superseded by the AI-assisted search that returns not just sources but a summary that purports to distill the relevant knowledge those sources contain. To people using this for the first time, it seems almost miraculous, and we can understand why many people (wrongly in our view) anthropomorphize the bot that can work such wonders. In our view, the ability to research a topic in this way makes teachers better; it follows that any teacher who wants to be at the top of his game should be using AI as a research tool. It will make him a better student of his own discipline.

This is especially helpful at a school like The Heights, where parents and administrators trust the teachers to exercise substantial discretion not only about how to teach but about what to teach. Many courses at The Heights consist entirely of content chosen by the teacher, so the teacher's ability to find the truth matters more than it would in a system where teachers must stick to a prescribed lesson plan for each class.

Moreover, in any discipline, great teachers are always seeking new ways to inspire the students themselves to want to know more. Sparking that inquisitive urge is much easier when teachers can use their own scholarly interest as a guide, asking and researching the same kinds of questions we want our students to ask. AI-assisted research can make teachers much better at this.

There are, to be sure, good reasons to be wary of the responses one receives to an AI-assisted research query. Sources may be slanted, or misunderstood, or in some (rare) cases wholly fabricated. In any important matter, one must read beyond the summaries and go to the sources themselves. But the same is true of any alternative method of research; the rule for scholarship must always be *caveat lector*.

In addition, there is some danger of dulling the knife; some danger that the teacher will (perhaps gradually) come to rely so heavily on the artificial assistance as to become dependent upon it. Newer teachers will probably be especially vulnerable to this danger because they are in a steeper section of the learning curve. A newer teacher must therefore be especially conscientious about courting some intellectual struggle, not only to avoid AI dependency but also to develop his own abilities more fully. But for most Heights teachers, after the first few years, the danger that AI will make a teacher lazy or distracted or burned out might not be in principle any greater than the risk that a teacher will become lazy or distracted or burned out without AI. In any event, it seems like a risk we can trust our teachers to manage on their own, as we trust them to strike a hundred other balances in the classroom.

B. Teachers as Curators of an Intellectual Tradition

We also see generative AI as an important asset for teachers as curators of the disciplines they teach. By curation, we mean everything from the construction of the syllabus (in which the teacher as curator must decide what the class will cover) to the formulation of the test questions (in which the teacher as curator must decide what the class *did* cover). Between these two endpoints, a teacher acts as curator when he chooses the readings or other instructional materials, when he chooses what to say in class about them, and when he chooses problem sets, essays, or other exercises to assign. In all these phases of curation, we have seen examples of how generative AI can make teachers much better curators of the knowledge they wish to present to their students.

AI can be particularly useful at the syllabus stage by suggesting additional readings to supplement or complement those already familiar to the teacher. All of us, no matter how widely read we are, can benefit from these suggestions much as we benefit from the suggestions of a human colleague. The suggestions will not all be good, but neither will those of the human colleague. We learn from bad ideas as well as good ones, as long as we insist on exercising the very human faculty of discerning what is good and true.

Teachers can also use AI to generate a far greater number of problem sets or test questions based directly on those the teacher is already using. This is one of the ways in which we might imagine AI as a shortcut for the teacher, making his job easier. But as a practical matter, that is probably backwards. Probably, the teacher's effort remains constant and the number of problems or test questions expands. That at least has been the experience of many faculty members who would otherwise generate fewer problem sets, or reuse more of the questions they have. In addition, some Heights teachers are using generative AI to create *remedial* problem sets for students to do after they've performed poorly on a test. This is not (yet) the techno-optimists' dream of an AI tutor for every student. But for the teacher who knows he has taught this material already in a way that has led the vast majority of students to mastery, the possibility of using an AI bot to create these remedial resources by mimicking and adapting the exercises the teacher has already created is a significant effort multiplier.

Finally, generative AI can help us present our lessons to students in a more compelling way. Most obviously, AI can create charts, graphs, or maps that make exactly the point we want to make in class, which is indispensable when a picture really is worth a thousand words. In addition, generative AI can critique and improve our lecture notes, our discussion questions, and our essay prompts. All of this works to the benefit of our students.

We hope not to oversell this "presentation" advantage, and so we wish to distinguish it immediately from some of the flashier applications of AI that are currently on the market. For example, starting with an upload of the assigned readings and the instructor's lecture notes, Google's Notebook LM platform can create a podcast in minutes in which two artificial voices simulate a fairly lengthy review of all the material in the course. This is amazing, but probably not useful because no student would take it seriously *except* as a shortcut to avoid doing the real work of engaging with the material.

From the teacher's point of view, there is here again some danger of rendering the teacher dependent through over-reliance on AI. AI researchers have already documented a [homogenizing effect](#) that tends to result from greater use of AI, and this would undercut one of the real strengths of our curriculum at The Heights, which otherwise benefits a great deal from the principle of subsidiarity. At the extreme, the teacher as curator might simply disappear; the course content might be found by AI, selected and edited by AI, and discussed in class according to questions drafted by AI. This would be catastrophic, like teaching from a textbook written and edited by the richest people in Silicon Valley. But again, on balance, we think we can trust our teachers to avoid these bad outcomes without forgoing the demonstrated benefits of AI use by teachers. That is why we have good teachers here instead of bad ones.

C. Teachers as Assessors of Intellectual Achievement

Some Heights teachers have already found generative AI useful even in the evaluation (grading and criticism) of student work. This is an area in which we urge caution, because our students rarely do their best work when they think we are not watching. However, assuming diligence and prudence on

the part of the teacher, we are convinced that generative AI can push us to provide better feedback and more of it, to the students' benefit.

A very pedestrian example is to ask AI to convert handwritten work to typed output for the teacher to grade. This has the advantage of allowing teachers to accept handwritten work (thus ensuring that the work is more likely the student's own) without forcing the teacher to decipher each student's writing (which is in many cases a challenge if not a penance). This is especially useful for teachers who use blind grading, because some students write so distinctively that it is impossible not to know whose essay is being graded. It may also remove at least one disincentive to assigning more written work that requires the teacher's personal attention to grade.

Going one step further, some teachers have asked AI to provide a second opinion on students' written work, and they have found that AI output sometimes points them toward truly commendable aspects of student work that they missed on a first pass through the work. This kind of second opinion might only confirm the teacher's own judgment, which is at least a minor benefit; but it might also push the teacher to "think twice" about differing evaluations of student work. In all cases, our experience has been that it results in more and better personal engagement with the students' own efforts and better feedback to students.

One English teacher regularly uploads student essays to Claude and asks Claude for a critique and a grade, which he then compares *in class* to his own critique and the grade he assigned. Students come away having learned not only a thing or two about what makes a good essay, but also some important generalizations about all feedback: that it is often right, never infallible, but usually helpful even when wrong because it can stimulate us to reaffirm what is right.

Finally, generative AI can help teachers identify patterns in student work that can inform the teacher's own efforts to improve in-class pedagogy. Which questions on the last test did most students miss? Was there a common theme that the teacher expected to see in the essays but did not find there? What were the most common grammatical mistakes or misspellings? Which foundational concepts were the worst-performing students still lacking, thereby disabling themselves from future advancement? A teacher *could* do any of this without AI, but most of us don't, or at least we don't do it as often or as systematically as we should. Most of us don't have the time.

In all these cases, as with research and curation, the risk of a teacher's overuse persists. But with grading, there is the additional risk that students might come to perceive the grade as having come from AI rather than from the teacher. It is imperative that teachers prevent this perception from arising. We recommend not only unflagging diligence from the teachers to make sure they are still the primary graders, but also complete transparency with the students so the teacher-student relationship remains strong and unmediated by machine.

And here we have one solid disincentive to overuse by teachers: task force members who have pushed themselves to try these new tools has found that they have made us *busier* and *more ambitious* about what we can give the students. Handwritten essays that have been transcribed in easily written type get read more closely; grades and critiques that are subjected to a artificial "second opinion" acquire more depth and more nuance. This is all to the students' benefit.

III. AI Presents More Peril than Promise for Students

The promise of AI for teachers consists primarily in its ability to improve our output—that is, to make our teaching better so that students receive better instruction in all of its dimensions. The corresponding risks in each situation consist primarily in the danger that teachers might become weaker in “process” terms; that AI will dull our intellects and reduce our effectiveness in human terms with our students. We believe the focus on output is the right focus for teachers, who all possess the requisite capacity and formation to do their jobs well without AI.

But with students, the focus must be on the development of the students’ abilities rather than the quality of their output. Our goal for our students is not *simply* for them to produce better work; it is for them to develop the intellectual skills necessary to produce better work. For students, process matters more than output. This inversion of the relative value we place on substance and procedure leads us to weigh the risks and opportunities very differently for AI use by students than we do for teachers. In short, we believe there should be a very strong presumption against even the *optional* use of AI by students, and we can imagine almost no good reason ever to assign work that *requires* students to engage with generative AI.

In most situations, asking a student to do his homework with the help of generative AI makes about as much sense to us as a coach buying expensive hydraulic lifts so that his players can do more pushups: it defeats the purpose of assigning the work at all. The difficulty of our assignments is not a bug in the educational process; it’s a feature. And we don’t actually need better essays from our students on *The Federalist Papers* or *Les Misérables*—we can find plenty of good essays to read already, or we can write them ourselves. We aim to create not better essays, but better essayists; not better scholarship, but better scholars.

It follows that the harm caused by AI depends on the developmental capacity that it arrests or diminishes. Thus, younger children, who are developing the most fundamental cognitive abilities, are hurt the most by being “helped” to artificial competence in these same areas, because their diminished capacities will hold them back later. This should make us much more reluctant to permit *any* exposure to AI platforms in the younger grades, because it is almost impossible to identify anything the AI platform does for the student that he would not more profitably do for himself without AI. In the upper grades, we begin to see some possibilities for our *oldest* and *ablest* students to benefit from *limited* exposure to AI platforms, but all three italicized words mark important limits.

One capacity we do *not* seek to develop pedagogically is often described as “AI literacy.” The idea is that AI isn’t going away, so students need to know what it is and know how to use it well rather than poorly. We are skeptical that anything about AI that truly needs to be taught is worth teaching to students between the ages of 8 and 18. It is true, in a sense, that AI is not going to go away; but it is also true that any particular AI interface that we could use with students today will be receding in our rear-view mirrors by the time we complete the “AI literacy” curriculum. What students really need to know about AI is what they need to know about nearly every other source of information: question it; scrutinize it; decide whether it’s true. In that sense, becoming a competent user of AI in the future depends not on what one learns about AI, but on what one learns about the whole rest of the universe, *excepting* AI.

To put these principles into practice, we propose two hard and fast rules that should apply to every student at The Heights, regardless of age or ability:

1. **No student should ever, for any reason, let AI read anything the teacher has assigned for the student to read.**
2. **No student should ever, for any reason, let AI write anything the teacher has assigned for the student to write.**

Violation of either of these rules invariably results in a loss of agency for the student, regardless of the student's age or ability, and therefore works directly against our mission in all cases. More generally, we believe there should be a fairly strong presumption against *any* AI use by students, though many of us think there might be limited circumstances in which it can benefit our older and abler students.

The consequences of the temptation to shortcutting that generative AI inflicts upon our students are familiar to anyone in contact with current college students and recent graduates. Most members of the Task Force, and presumably most members of the Administrative Council and the rest of the faculty, will have heard from our own children and our recent Heights alumni about the complete collapse of academic rigor that has been occurring throughout higher education. We have heard of a student struggling to do his own very demanding work who was told by his professor that he should just get AI to do it for him, like everybody else. Princeton has decided to abandon its honor code and require in-person proctoring of exams for the first time in over a century. Temptation is the right framework for this, not unlike other online temptations that prove too powerful for many students to resist.

And without question, a student's reliance on AI during the formative years of education will diminish that student's professional effectiveness later. A mid-level lawyer in one of DC's leading law firms recently reported to one of us that it seemed that only a few of the junior lawyers at her firm had really been to law school; the rest, she suspected, had asked AI to read all the cases for them—and were now asking AI to write all their briefs, which pre-empted the further development of their skills that should have been occurring in those years (generally their late twenties!). This is the problem with student use of artificial intelligence: it may work well enough for the student in the artificially constructed world of homework assignments and semester grades, but it fails to develop real intelligence in the student, and real intelligence is what the student will need in the real world.

For this reason, we think virtually any student interaction with AI should be discouraged, even if the interaction is strictly optional. One possible exception that we discussed repeatedly is the old-fashioned research paper. An 11th- or 12th-grade student who is researching a long thesis can benefit from AI research in much the same way that a teacher can benefit from AI-assisted research. Indeed, anyone who uses Google for web searches will find it difficult to avoid this, as Google now reports search results with an AI-generated summary (by Gemini—you've been warned). In this context, the net effect of the student's use of AI as a research assistant is *arguably* to support deeper investigation and therefore to increase academic rigor. But as ed-tech writer Andrew Cantarutti [points out](#), the ability to discern the relevant sources from among the irrelevant dross is itself one of the capacities we wish to develop in the student, and a student who turns to AI for research will not be developing or exercising that power of discernment as rigorously. The AI giveth, and the AI taketh away.

Another way to give students practice at working with sources without exposing them to AI would be for the *teacher* to consult AI to create a *packet* of resources for further research, which could then be given to all students so they can develop the ability to read syntopically across various texts and generalize from all of them, without the danger that AI will cut their meat for them. But it is perhaps only fair to note that in this case the *teacher* has cut their meat for them, which is perhaps better for the student-teacher relationship but will still not give the student practice at rejecting the irrelevant

(unless the teacher stocks the research packet with “decoy” sources, which seems like an evil genius maneuver). There were trade-offs in teaching long before there was AI.

What about a student soliciting editorial *help* from AI with that student’s written work? For example, a student might upload an already-written essay to a generative AI platform and ask for suggested improvements of the kind that have traditionally been provided by parents for at-home assignments. Would it be better for the parent to provide the suggestions personally? Probably, depending on the parent’s own strengths or weaknesses in this area. Would it be better for the student to seek feedback from the teacher? Certainly, but the teacher cannot be instantly available to the student at all times to provide individualized feedback, whereas AI platforms can. Theoretically, this could benefit our oldest and ablest students; the strongest case would be an excellent student for whom English is a second language not spoken in the home, and who has the character to take constructive suggestions from an AI platform as an opportunity for *more work* rather than less. Practically, the number of students we have just described is quite small.

Whatever one thinks of these edge cases, we generally think teachers should structure their assignments and assessments to minimize any opportunity for students to use AI. That might mean that more of our writing assignments will shift from at-home composition to in-class composition. It might mean that we should prefer handwritten work from home to typed work from home. It might mean more reliance on oral exams in which the thinness of AI-assisted learning will show itself. It might make us more wary in general of the weight we assign to the homework grades, lest we inadvertently increase the student’s temptation to use the AI shortcut. Keeping the students away from AI and inducing them to do their own work should be each teacher’s responsibility as well as each student’s.

Because we place such high importance on the prevention of AI shortcutting by our students, we have also investigated various AI-based services designed to detect illicit use of AI. So far, our verdict is that these products are not very effective. Furthermore, these are not covert detection systems; they all operate in plain view of the student, and we think that communicates a mistrust that might damage the student-teacher relationship and make dishonesty a self-fulfilling prophecy. We reluctantly conclude that while the *desired* level of cheating is zero, the *optimal* level of cheating is not zero because the detection efforts that would be required to drive cheating to zero all have disadvantages that are worse than the cheating.

We propose language for the student handbook below, but this way of living with generative AI should not be left to a few words in a handbook. Teachers should instruct their students that any use of AI to read assigned texts, or generate the required student work product (including not only essay-writing but also problem-solving in math or science and exercise-completion in any subject) is a form of academic dishonesty that both degrades and deforms the person who practices it. It is a crime that mostly victimizes the criminal. This message must be made crystal clear to students throughout the school if we expect it to become part of our school’s culture. We recommend the use of some of the existing monthly assemblies for this purpose, and we encourage the core teachers to take the lead in creating and sustaining a healthy culture that rewards our students for wrestling manfully with hard questions instead of settling indolently for easy answers.

The firm prohibition we recommend on student use of AI necessarily creates line-drawing problems. One set of problems arises from the question of how much AI is too much AI: Is it licit to ask for suggestions for an essay as long as the student does not accept them *in haec verba*? Is it licit to ask for three or four examples from a primary text that might serve as evidence for a thesis of the student’s own invention? Is it licit to ask for three or four suggested theses? We do not presume to

resolve all of these cases for every assignment in every class. We are confident, however, that “don’t use AI for your assignments” will be the right message for the vast majority of our students in the vast majority of cases. Each teacher can adjust the message appropriately to reflect his own expectations.

The other line-drawing problem arises from the way that generative AI is increasingly embedded in electronic tools we often need in order to do the exact type of work with which AI most directly interferes: writing. It is good for students to type at least some papers (and also good for them to write some papers in class), so the temptation will be strong for the student to have AI open at the same time. Similarly, there are some kinds of specialized STEM content online that we think it is good for students to view; the temptation will meet them there as well. Whatever teachers decide about these cases, it is important—probably more important than usual—for teachers to communicate their expectations explicitly to parents as well. (A laudable example from Michael Hude is attached as Appendix A.) This is implied not only by our weighing of the pedagogical tradeoffs, but by our well-known philosophical stance against the disruptive effects of screens on the entire educational enterprise, which parents will (mostly gratefully) take as the default expectation with regard to AI. No parent should have to wonder whether Johnny really has to go online to complete his homework; with very few exceptions, it should be clear that anything found online is entirely optional.

Parents will of course make their own decisions about how much they will permit their children to interact with AI chatbots, and in what contexts. But familiarity with these tools will vary widely across the parent community, and we should expect many parents to look to the school for guidance on questions like which service to use, what limits can and should be imposed on use by adolescents, and whether there are any worthwhile technical barriers or guardrails they should implement at home. These practical questions deserve more consideration and experimentation than we have been able to give them up to this point, and given the pace of innovation in generative AI it may be unrealistic to expect any answer to remain true and complete for very long. It may be, however, that a working group should be formed to test various technologies and share our findings periodically with parents via *The Heights Forum*, in much the same way that we have shared information about alternatives to smartphones.

IV. Additional Questions About Technology and Pedagogy

Our Task Force on AI has repeatedly bumped up against related questions about technology and pedagogy that we recognized as relevant but which we considered beyond our charter. We include a few words here about these issues, in the hope that other groups of colleagues will continue to address them.

A. Writing in the Curriculum

Writing instruction was in some sense the canary in the coal mine with respect to generative AI, because students and teachers alike saw immediately that the technology could be used to commit almost undetectable academic dishonesty. We know from personal experience that many teachers (especially in humanities *other* than English) did in fact change the way they approached writing assignments; there may not have been any conscious intention to assign less writing, but that seems to have been the result.

The technologists who are most bullish on AI might well respond to this with a shrug. “So what if students are writing less? They won’t need to be very good writers, because AI will write for them.” We could not disagree more. As we have noted already, reading and writing seem to us the most quintessentially human skills we teach our students, and we should go to great lengths to protect them.

Regrettably, this is once again an area in which it is our average students who are probably most at risk. We continue to graduate many truly fine writers, and we presume that these fine writers would be the students most likely to resist the temptation to use AI for their writing, because they evidently love writing and are good at it. But our aim should be to give *all* our students the ability to express themselves clearly and persuasively. This is largely a matter of writing and re-writing often, and of course this does the most for the students who find it most difficult. But it is precisely those who find writing the most difficult who will be most strongly tempted to use the AI shortcut.

We believe teachers can mitigate the risks of academic dishonesty by structuring their writing assignments somewhat differently, but there are significant trade-offs. If we are assigning more in-class writing and fewer at-home assignments, we will probably get less cheating, but we will probably also get less writing and almost certainly less long-form writing. Is that a good trade?

The conversations necessary to weigh the pros and cons of particular adaptations in our writing instruction must include other teachers sitting around other tables. But we think AI has forced upon us the need to re-examine our approach to writing instruction to make sure we continue to do it well.

B. Homework, Online and Otherwise

While our mandate was focused on the rise of generative AI and what it implies for our pedagogy at The Heights, we inevitably found ourselves talking about other computerized learning resources, if only because educational technologists typically imagine AI existing alongside these resources rather than superseding them. There are, without question, many great online resources out there, and the number is almost certainly growing. Math Academy is a standout in this domain, and some of the online resources provided by the College Board for our AP classes are also very good.

But at the same time, *personal* instruction remains the cornerstone of our educational approach, and research reported by Dr. Jared Cooney Horvath in his book *The Digital Delusion* gives us several important reasons to prefer it precisely because it is personal. Students are more attentive to teachers than they are to screens; students are more motivated to learn from teachers than from screens; and students are better able to transfer the skills they have learned from one domain to another when they are guided by a human teacher. Thus, despite our admiration for some of the computer-based resources currently available, we think there is always a cost to shifting the student’s focus from a trusted teacher inside the building to a screen projected by a team of strangers outside the building.

The situation is probably better with respect to resources created by the teacher himself, which is currently the case for the optional logic and Latin resources created and posted by the same Heights teachers that students encounter in the classroom. But even these optional resources still create the potential for fairly long stretches of engagement with computer screens at home. That makes it harder for our community of parents to set and enforce sensible rules about screen use at home.

The Heights has been ahead of the curve in protecting our students’ attention spans from the balkanizing influence of smartphones and social media. In light of the harmful effects of mediating screens in general, and the even more harmful effects that AI can have on student agency, we have had to consider what role these computerized resources should play in a Heights education. Our answer in a nutshell is: it depends.

In our view, it depends first of all on what we assume the alternative to be. Watching AP science videos online for homework is very uncharacteristic for a Heights homework assignment, but in a high-level STEM course this may be justified. By contrast, directing students to an online dictionary of Latin or Greek is probably *not* justified, because a book would do the job almost as well and would probably result in much greater long-term retention. The logic lectures currently online are there not to replace or even supplement what's offered in class, but to make what's offered in class available to a student who misses class and would otherwise be limited to what is in the notes he borrows from a friend.

We also think it is easier to justify *permitting* the use of online resources than to justify *requiring* such use—but even here there are pitfalls to avoid. Is the online resource so helpful that not using it would put some students at a significant disadvantage? This sets up a true dilemma, because if the online resource is *not* so useful then it's hard to see why we should offer it at all; and yet if the online resource *is* really useful then it's hard to see why we don't assign it to all students, to bring them all up to the desired level of mastery.

Finally, there is the question of the extent to which online assignments affect our students' grades. Routine homework assignments are commonly graded for completion rather than perfection, precisely because we understand that the student's encounter with the material helps to produce the desired learning outcome even if some of the answers are wrong. But unfortunately, this is an area (like at-home writing) in which we think the availability of AI has altered the balance of risks and benefits. In math, for example, assigning a completion grade for each problem set completed at home invites students to take a picture of the problem set and ask AI to produce the answers. We are certainly not saying that teachers should stop assigning homework grades; we are saying only that the more the homework "counts," the more some students will be tempted to outsource it to AI. And one of the most effective ways to protect students from the harmful effects of taking the AI shortcut is to create assignments that need not and perhaps cannot be completed with the aid of a computer.

We do not imagine we have given these problems a thorough consideration here, and we do not presume to resolve all these questions for every teacher. But we do think directing a student to an online resource is unjustified unless (1) the benefits are demonstrable rather than speculative, and significant rather than negligible; and (2) there is no suitable offline alternative. Again, there will be exceptions—probably many exceptions in this case—and this is a matter that must be committed to the teacher's sound professional judgment, in consultation with the head of his school.

V. Proposal for a Heights Learning Lab

We spent some of our early sessions investigating the extravagant claims made for AI-based education by Alpha School, a well-funded chain of private schools that claims to educate students more than twice as fast as a regular school can in just two hours per day. How? By replacing teachers with AI. But upon closer examination, Alpha seems to be using *computer*-based resources, but not *AI*-based resources; for example, it uses Math Academy for math instruction and Duolingo for language instruction. Math Academy is great, but it's not an AI-based personal tutor. More fundamentally, the leaders of Alpha School seem to have an impoverished view of education: perhaps as "the efficient delivery of measurable amounts of information to students who are properly incentivized to receive it and give proof of receipt." Why? They don't say. It seems to be school for school's sake, uninformed

by any conception of final cause. No one on the Task Force thinks The Heights should be moving in that direction.

And yet: in any discipline, there are certain foundational skills that a student must master before he can progress further. We can think of the multiplication table in math, for example; or in English, basic spelling and grammar, or the ability to write a well-conceived paragraph. The failure to acquire these skills doesn't just result in a bad grade on one test or one paper; it holds students back indefinitely, unless and until the material is finally learned. Fortunately, these skills can typically be acquired through structured repetition—drill. Unfortunately, when students fail to learn these building-block skills when they should, the train that carries most of the students forward often leaves the station without the straggling few who failed to attain mastery with their peers. In cumulative courses like math or Latin or logic, the number of stragglers can grow steadily throughout the year as more and more people fall off the back of the train.

This is where something like Alpha's AI-based model would be useful, if it existed, because it could ask the student questions, give the student hints, correct the student's answers, and ask the student more questions. It sounds too good to be true, and we think it is; but that does not mean that *if* it became true it would not be good. And indeed, even the computer-based resources that already do exist [can be a big plus for students who would otherwise be left behind](#). Research on the effectiveness of such "adaptive digital tutors" suggests that they are effective when used within narrow domains of knowledge with clear right-or-wrong answers, like math or Latin or logic. Some research suggests they are [even more effective when a human teacher remains "in the loop,"](#) with the human+AI team besting either AI or a human tutor alone. Could AI help us provide this kind of targeted remediation, at the student's own pace and on the student's own time? We think it might, and that the possibility is worth pursuing further. We therefore recommend that the Administrative Council approve and fund an initiative we have provisionally called the Heights Learning Lab.

As we envision it, the Heights Learning Lab would be a physical place, staffed by a teacher or IT administrator, where a student could use a school computer to access learning modules selected by his teacher to remediate knowledge or skill gaps the teacher has observed from the student's classroom performance. The modules would provide computerized drilling, but they would not really be AI-based resources; they would just be drills devised by the teacher. Where does the AI come in? The AI makes these fairly easy for the teacher to create, and based on some very preliminary experimentation we think they might be valuable.

We tested the model this past semester in the freshman logic classes—except that in order to get good data, we used the modules for test *preparation* rather than for remediation. There was time for four "phases" of testing.

- Before one chapter test, two of the four sections got the standard chapter review, while the other two sections (the two whose averages on previous tests had been *significantly* lower all year) used a computerized study module constructed by Mark Grannis with the help of Anthropic's Claude service. (The module has since been refined, but you can get a good sense of it [here](#).) Grannis circulated among the desks while the students answered the review questions, sometimes offering hints and sometimes explaining the answers the students didn't understand. For the first time all year, the two sections who had the computerized prep outperformed the other two.
- Then, for the next chapter test, Grannis gave the two first-period sections [this](#) computerized review while giving the two third-period sections the normal review. The first-period averages had been virtually identical with the third-period averages all year, but

for this chapter test the students with the computerized review outperformed the students with the in-person review from Grannis.

- After both of these chapter tests, students who had bombed the tests were invited (optionally) to get extra credit by completing remedial modules at special “Lunchtime Logic” sessions. (Examples of the remedial modules are [here](#) and [here](#).) There was not time to offer a “retest” to see how well these students filled their gaps, but the modules were widely used (perhaps in part because of the novelty).
- For the final chapter test, all students received [this](#) computerized review, and were surveyed about the experience. For what it’s worth, the students reported that they thought they learned more from the structured repetition of the computerized review. Grannis’s observation is that the review modules, by enabling each student to work at his own pace during the session, led to significantly more practice by all students, and perhaps especially those in the middle of the class academically.

This is not what any good social scientist would accept as good data, but it seems to us a promising avenue for further development.

For some subjects, resources of this type already exist; in particular, Math Academy is one we have mentioned with approval several times already. But even though the logic resources had to be created from scratch, this was not actually very difficult. The teacher simply has to identify the skill or competency that requires remediation, construct some exercises that address that skill or competency, and then work with Claude to create and refine the code. Tom Cox, George Martin, and David Maxham are currently working on similar modules for Latin. Some of us think this could even be useful for addressing deficits in reading comprehension, using short excerpts such as those that are used for the Classical Learning Test; others of us don’t.

The biggest obstacle we see, other than the teacher’s time and the almost nominal cost of a souped-up AI subscription, is that we strongly believe the Heights Learning Lab should be a *place*, staffed by a *person*. Or to put it another way, we do not believe students should be sent online for these resources from home. Rather, the student who has been sent to the Learning Lab should know viscerally that he is in school doing school work. We understand that classrooms are scarce, but we think this is worth trying.

VI. Formal Recommendations of the Task Force

A. Revision of the Handbook Policy on Academic Honesty

Under the heading of “Academic Honesty,” the Handbook rightly frames academic dishonesty as an injustice against others that is ultimately damaging to the dishonest student, “by replacing honest effort with fraud.” The handbook currently lists four types of offenses against academic honesty, including “2. plagiarism or the unacknowledged borrowing of the work of an author;” and “3. using unauthorized aids”.... We believe this is where the prohibitions on illicit use of AI should be placed. Accordingly, we would revise that paragraph to read as follows:

A student may commit an offense against academic honesty by:

1. cheating on a quiz or test;

2. copying another student's homework;
3. plagiarizing, or presenting as one's own the work of another author, including any unattributed or misattributed use of text generated by an AI platform or product; or
4. using unauthorized aids, including the use of an AI platform or product to summarize, excerpt, or digest any assigned material or any work the student represents himself as having read.

Conclusion

We began by observing how quickly the current challenges posed by AI came upon us. We end by observing that we have not seen the last of these challenging developments and the pace of change is likely to be brisk for the foreseeable future. We have learned much from each other during the last six months: what AI can do, what teachers can do (and are already doing) with AI, and what students can do (and are already doing) with AI. We have also benefited personally from the opportunity to reflect on what teachers do that AI *cannot* do: relate to human students and help them grow in virtue, understanding, and genuine insight. We hope the Administrative Council has found this Report worthwhile, but we are confident we have not offered the last word on AI at The Heights, and we hope the Administrative Council will make this a subject of continued study and faculty formation.

Appendix A: An Exemplary E-Mail to Parents on AI Assistance at Home

Dear Parents,

I hope you and your family had a nice Easter vacation!

So far this year we've done mostly in-class essay writing, which is an increasingly important skill, as schools and colleges have transitioned to this format (mainly for tests) to guard against the use of AI-assistance.

Writing outside of class, though, is still a very important skill for students to sharpen, and so as we finish the last quarter of the year, I'll be assigning more of these sorts of essays. The ability to take one's time to articulate one's thoughts clearly and artfully—and without as much of a time restriction—is likely where the most growth takes place when it comes to writing.

But the danger of students seeking outside help is elevated with so many resources at their disposal. **And so I am asking you to partner with me to do what we can to minimize these risks. Please be on the lookout for your son seeking these sorts of aids.**

While I consider seeking AI help to be impermissible, I am firmly in support of your son sitting down with you or a peer to go through the editing process. We'll in fact be doing in-class peer editing quite a bit in the next few weeks. But having a more mature presence at home go over their writing with them is probably more valuable. Thus I strongly encourage you to do that with your sons.

Thank you for all your support this year; it's been a true joy to work with your sons! I am very grateful for the time I have had with them and hope to make the most of the remaining time we have together.

Please be in touch with me if you have any questions or feedback.

Many thanks,

Michael R. Hude