

Virtue Ethics and Student Use of Large Language Models in Engineering Education

1st Zachary Deutsch
Department of Mechanical Engineering
Duke University
Durham, USA
zachary.deutsch@duke.edu

2nd Rich Eva
Pratt School of Engineering
Duke University
Durham, USA
richard.eva@duke.edu

3rd Siobhan Oca
Department of Mechanical Engineering
Duke University
Durham, USA
siobhan.rigby@duke.edu

Abstract—This paper examines the ethical implications of student use of large language models in engineering education through the lens of virtue ethics, with a focus on practical wisdom, truthfulness, and temperance. As students increasingly use large language models for writing, coding, and problem solving, it is important to understand their impact on learning. Current discussions focus on academic misconduct and authorship, but provide limited insight into how large language models influence the development of intellectual and moral virtues that are important to engineering education. This paper examines these virtues and applies them to representative scenarios of student use of large language models, using literature from engineering education and artificial intelligence ethics. This paper argues that large language model use may support or undermine development of practical wisdom, truthfulness, and temperance depending on whether it replaces intellectual engagement or supports it while preserving student responsibility for deliberation, judgment, and understanding.

Index Terms—engineering ethics, engineering education, artificial intelligence, virtue ethics

I. INTRODUCTION

Students are increasingly using large language models (LLMs) for writing, studying, problem solving, and other academic tasks [1]. Their availability has prompted concern about how they affect student learning, specifically in engineering education [2]. According to the Accreditation Board for Engineering and Technology, engineering education should develop the ability of students to identify, formulate, and solve complex engineering problems, apply engineering design to produce solutions that meet specified needs, and make ethical judgments [3]. We propose an ethical framework to understand how students should use LLMs. This paper builds on previous work examining engineering students' interactions with LLMs and extends that discussion by analyzing student LLM use specifically through the Aristotelian virtues of practical wisdom, truthfulness, and temperance [1], [2], [4], [5]. Practical wisdom concerns whether the use of LLMs preserves student deliberation and judgment. Truthfulness concerns whether student work genuinely reflects their own understanding. Temperance concerns whether student use of LLMs is moderate rather than dependent. The paper defines these virtues, examines LLMs in engineering education, and evaluates their implications for students' intellectual and moral development.

II. VIRTUE ETHICS

Aristotle defines virtue as a disposition that guides reason and action. In *Nicomachean Ethics*, he writes that virtue is “a state concerned with choice, lying in a mean relative to us, this being determined by reason and in the way that the person of practical wisdom would determine it” [5]. Aristotle distinguishes between intellectual virtues and moral virtues.

Moral virtues are character traits that precede emotions and actions. Aristotle states that “moral virtue comes about as a result of habit” [5]. Individuals therefore become virtuous by regularly performing virtuous actions until these actions become habits. Intellectual virtues, on the other hand, are about finding truth through reasoning. Aristotle explains that “intellectual virtue in the main owes both its birth and its growth to teaching, for which reason it requires experience and time” [5]. These two types of virtue are closely related because morally good action requires sufficient reasoning.

A. Practical Wisdom

Practical wisdom is a type of intellectual virtue. Aristotle defines practical wisdom as “a true and reasoned state of capacity to act with regard to the things that are good or bad for man” [5]. Because practical wisdom is about action in particular situations, it requires experience and sound judgment about appropriate means, both of which develop over time rather than from abstract principles alone. Practical wisdom therefore informs deliberation about how to act, allowing one to determine appropriate means toward desired ends.

The relationship between practical wisdom and action can be thought of as decision-making, or deliberation and choice [5]. Deliberation is the process by which one reasons about possible courses of action and considers the possible means to achieve one's desired ends. Choice is the decision about a course of action following deliberation.

For clarity, this paper distinguishes between several related concepts. *Reasoning* refers broadly to the cognitive process of analyzing information and drawing conclusions. *Deliberation* is a specific form of reasoning that concerns evaluating possible courses of action and their means. *Choice* is the selection of a particular action following deliberation. *Decision-making* refers to the overall process that includes both deliberation and choice. *Judgment* refers to the quality of these processes,

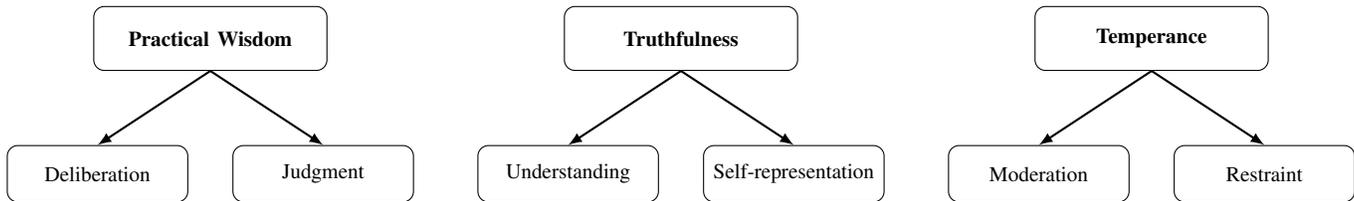


Fig. 1. The three Aristotelian virtues used to evaluate student use of LLMs in engineering education and the central concepts associated with each virtue.

including how well one evaluates alternatives and selects appropriate actions. *Practical wisdom* is the intellectual virtue that governs these processes, allowing for sound deliberation and judgment [5].

B. Truthfulness

Truthfulness is a type of moral virtue. Aristotle describes truthfulness as the virtue of a person who is “truthful both in life and in speech, because he is so in his state of character” [5]. He describes it as a mean between boastfulness, which exaggerates the truth, and self-deprecation, which understates it [5]. Truthfulness therefore involves representing oneself as one is: neither saying more nor less than is true about oneself. Truthfulness is important because it reflects a stable disposition toward honesty.

The relationship between truthfulness and action can be understood through genuine representation. In speech and behavior, individuals present their character to others. Truthfulness guides this presentation by requiring that it correspond to what one’s character actually is. Truthfulness helps to ensure that one’s words and actions are aligned [5].

C. Temperance

Temperance is a type of moral virtue. Aristotle defines temperance as a “mean with regard to pleasures” [5]. Temperance regulates desire by guiding people to take pleasure in the right things, in the right way, and to the right extent. It is therefore a virtue of moderation.

A temperate person is not free from desire, but has learned to desire appropriately. Aristotle contrasts temperance with self-indulgence, which is the excessive pursuit of pleasure, and with insensibility, which is the unreasonable lack of response to pleasure [5]. Temperance thus represents a balanced disposition in which desire is regulated by reason.

The relationship between temperance and action can be understood through the regulation of desire [5]. In situations involving pleasure, people may be drawn toward what is immediately satisfying rather than what is best. Temperance supports good action by preventing pleasure from superseding sound judgment. It helps ensure that one’s actions are moderate, which is important when considering the various ways one might use an LLM.

III. LARGE LANGUAGE MODELS

LLMs are a type of neural network that can process and generate text. As LLMs are trained on increasingly large datasets with greater compute, they demonstrate the ability

to not only generate text but also perform various language-related tasks [6], [7].

Despite these capabilities, researchers have identified several limitations of LLMs. One is hallucination, in which LLMs generate outputs that appear plausible but are factually incorrect [8], [9]. Because LLMs generate text by predicting likely sequences, their outputs may be factually incorrect. Some argue that LLMs use statistical pattern recognition rather than actual understanding of language or reasoning [9]. For example, Searle’s Chinese Room Argument states that programming a computer may make it appear to understand language but could not produce real understanding [10]. Others argue that we need to reconsider the definition of understanding. In Daniel Dennett’s *The Intentional Stance*, understanding can be attributed to the entire computer system, not the individual parts [11]. Applied to LLMs, this perspective suggests that their apparent understanding can be attributed to their behavior as a whole, rather than to any internal representation of meaning. Nonetheless, because of the potential limitations of LLMs, there are ongoing discussions about the appropriate use of LLMs in domains that require sound understanding and judgment [6], [9].

IV. LARGE LANGUAGE MODELS IN EDUCATION

LLMs are increasingly used by students in education [1], [2]. One use is writing and studying. Students often use LLMs to write essay drafts, summarize readings, make study guides, and answer questions about course material. In this way, LLMs can function similarly to tutors, explaining complex ideas and helping students as they work through course material. Many students also use LLMs to brainstorm ideas and receive feedback on assignments before submitting them [1], [2].

LLMs are also used for technical coursework, including computer science, engineering, and math [8]. In computer science, for example, students often use and increasingly rely on LLMs for generating code, explaining algorithms, or debugging programs. In math, students working on problem sets can ask LLMs for step-by-step explanations of problems. LLMs are also widely used by students in engineering education to complete coursework.

V. ENGINEERING EDUCATION

Engineering education aims not only to teach technical knowledge but also to develop students’ ability for deliberation and judgment. Engineers regularly deliberate about design choices and make decisions about the potential consequences

of technology. The Accreditation Board for Engineering and Technology (ABET), in its *Criteria for Accrediting Engineering Programs*, outlines the ability to identify, formulate, and solve complex engineering problems, apply engineering design to produce solutions that meet specified needs, and make ethical judgments as important learning objectives for engineering programs [3]. Coursework in engineering education functions as evidence of what students understand and can do. The use of LLMs therefore raises not only questions about learning, but also questions about whether students’ work truthfully represents their understanding and ability.

VI. WHEN LARGE LANGUAGE MODELS SUPPORT OR UNDERMINE VIRTUE DEVELOPMENT

Whether LLMs support or undermine virtue development depends on how they are used. In some cases, LLMs scaffold student thinking while leaving deliberation and judgment with the student. In other cases, they replace these processes. The ethical significance of LLM use depends on whether it preserves intellectual and moral virtue. Across these virtues, the effects of LLM use can be understood through three use cases: an LLM may function as an assistant, a collaborator, or a substitute for the student’s work.

A. Practical Wisdom

Does LLM use support or undermine student deliberation and judgment? Judgment develops through active engagement with problems rather than through passive reliance that LLMs may promote [14]. The evaluation of LLM use therefore depends on whether students remain actively engaged in deliberation and judgment.

LLMs can support the development of practical wisdom when they scaffold deliberation and judgment. In this case, the LLM provides explanations or examples, while the student remains responsible for evaluating and applying this information. For example, a student working on a statics problem might ask the LLM to explain why one free-body diagram is more appropriate than another, then use that explanation in judging how to set up the equilibrium equations. Similarly, a student writing code might use the LLM to compare two possible debugging strategies, while still deliberating about which method best fits the structure of the program. In these cases, the LLM helps the student think through the problem without making the judgment for them. At the same time, this help still involves some tradeoff, since the LLM still offloads some deliberation by clarifying or organizing information for the student. The LLM contributes to the learning process, but it does not completely displace the student’s responsibility for deliberation and judgment.

In contrast, LLMs can undermine the development of practical wisdom when they replace deliberation and judgment. When students rely on LLMs to generate complete answers, they may bypass the processes of analysis and evaluation that are necessary to learning. For example, a student in a circuit course might paste a homework question into an LLM and use the full answer without deliberating about why the result

TABLE I
HOW DIFFERENT LARGE LANGUAGE MODEL ROLES MAY SUPPORT OR UNDERMINE VIRTUE DEVELOPMENT IN STUDENTS

Role	Potential effect on students
Assistant	May support learning by helping students evaluate and refine their work while preserving student responsibility for deliberation and judgment. Students remain actively engaged in the work.
Collaborator	May either support or undermine learning depending on whether students critically evaluate outputs and retain responsibility for revising the work. Its effect depends on how responsibility is distributed between the student and the large language model.
Substitute	May undermine learning by replacing substantial parts of student deliberation and judgment. This use may reduce engagement, distort students’ representation of their understanding and contribution, and encourage excessive reliance on the large language model.

makes sense physically. Practical wisdom must be developed through experience and engagement with specific problems [14]. In cases where students do not engage in deliberation, LLM use may reduce the development of practical wisdom.

There are also cases in which the use of LLMs is neither clearly supportive nor clearly undermining. When students use LLMs as collaborators, the LLM may take on some aspects of the deliberation process while the student retains others. For example, a student writing a lab report might ask the LLM to suggest possible explanations for an unexpected data trend, then compare those suggestions against the actual experiment and decide which interpretation is most plausible. In this situation, the impact on learning depends on how responsibility for deliberation and judgment is distributed. If students critically evaluate outputs, revise solutions, and make final decisions, then they may still engage in the processes necessary for developing practical wisdom. However, if students defer to the LLM’s outputs without sufficient evaluation, then the most important aspects of deliberation and judgment may be lost. The effects of LLMs depend on how they shape habits of thought, but this distinction is not always clear.

B. Truthfulness

Truthfulness concerns honest self-representation [5]. In education, this includes honest representation of one’s work contributions and understanding. LLMs can support truthfulness when they are used in ways that do not distort this representation. In such cases, the LLM may help students revise drafts or clarify ideas, while the student remains responsible for the substance of the work [1], [2]. If the student’s work still reflects their own contributions and understanding, then such LLM use does not necessarily interfere with truthfulness. The LLM assists the student, but it does not create a false representation of their contribution or understanding.

In contrast, LLMs can undermine truthfulness when they are used in ways that misrepresent the student’s contribution and understanding. When students submit LLM-generated or heavily LLM-shaped work as if it were entirely their own,

they present a level of understanding that they do not actually have. For example, a student might submit an LLM-written results section for a lab report even though they do not understand the results of the experiment. Or, a student in a programming course might submit LLM-generated code and receive credit for problem solving ability they did not demonstrate. Even if the final output is correct, the work may still give a false impression of independent thought. In such cases, LLM use may weaken the development of truthfulness by encouraging inaccurate self-presentation. It may also undermine the relationship between students and instructors. The ability of students to misrepresent themselves can shift the role of instructors from teaching students toward monitoring authenticity. In this way, the effects of untruthful LLM use may extend beyond the student's character to the broader structure of trust in the classroom.

There are also cases in which the effect on truthfulness is less clear. When students use LLMs as collaborators, the final work may reflect both student reasoning and LLM-generated contributions. For example, a student may outline an essay argument, ask the LLM for help reorganizing the structure, and then revise several paragraphs using phrases suggested by the LLM. In these cases, the ethical issue depends on whether the student genuinely evaluates and understands the work. If the student critically engages with the LLM's output and the final work still reflects their own reasoning in a meaningful way, then truthfulness, to some extent, may be preserved. However, it may not always be possible for the student to accurately recognize the extent of their own knowledge or contribution in these kinds of assessments. If the student presents the work in a way that exaggerates their own contribution or understanding, then truthfulness may be compromised.

C. Temperance

The central question for temperance is whether LLM use remains moderate or becomes excessive. Temperance concerns the regulation of desire and the avoidance of excess [5]. It is also shaped through habituation. In the context of education, this means asking not only whether students use LLMs in ways that support learning, but also whether repeated use creates a disposition to avoid difficulty or uncertainty.

LLMs can support temperance when they are used with restraint. In such cases, the student uses the LLM for limited assistance, such as clarifying a concept, checking an approach, or exploring alternative ways to understand a problem. For example, a student might ask an LLM for a simpler explanation of a math concept after first attempting the assigned reading, or use it to check whether their interpretation of a graph is reasonable before finalizing their own answer. The LLM serves only as a support for learning rather than a substitute for it.

In contrast, LLMs can undermine temperance when students use them with excessive reliance. Students may use LLMs not only for assistance, but also for avoidance of difficulty or uncertainty. In engineering education, this is especially significant because students are expected to develop persistence in the face of difficulty and confidence in independent problem

solving. Excessive reliance on LLMs may therefore weaken not only restraint, but also the habits that engineering education seeks to form. When students repeatedly turn to LLMs to solve problems that they would otherwise be expected to do themselves, their use of the LLM may reflect a failure of restraint. The deeper issue is not simply frequency of use, but the habit such use forms. A student who consistently turns to the LLM at the first sign of uncertainty may become less willing to persist through effort or develop confidence in independent problem solving. Although the effects of becoming excessively reliant on LLMs may not be obvious in the context of the course, it may become more consequential in engineering practice, where problems are less structured and require independent judgment under uncertainty. In such cases, LLM use may weaken temperance by cultivating dependence rather than restraint.

There are also cases in which the effect on temperance depends on the degree to which students use the LLM. Collaborative use may still be moderate if the student uses the LLM with restraint. For example, a student may use the LLM once near the end of an assignment to compare their reasoning with an alternative explanation, while still doing the main work alone. However, such use may become intemperate if reliance on the LLM becomes habitual or if the student increasingly turns to it whenever learning becomes difficult.

VII. HOW EDUCATORS SHOULD RESPOND TO STUDENT USE OF LARGE LANGUAGE MODELS

Engineering education must preserve student engagement in decision-making. Judgment cannot be developed through passive reliance on LLMs, but instead requires active engagement with problems [14]. LLM use should be structured so students remain responsible for deliberation and judgment. Accreditation standards emphasize that engineering education programs must develop students' abilities in problem solving, ethical reasoning, and responsibility [3]. These outcomes require students to deliberate about design alternatives, exercise judgment about trade-offs, and make decisions while considering the consequences of their work. Engineering education should therefore ensure that LLM use does not displace these processes.

One implication is that assignments and learning activities should be designed to require active engagement. Tasks should require students to explain their reasoning and justify their decisions. These forms of activity correspond to higher levels of cognitive development in Bloom's taxonomy, such as analysis and evaluation [15]. Such engagement is necessary for developing judgment [14]. By structuring assignments in this way, instructors can reduce the likelihood that students rely on LLMs as substitutes for their own work.

A second implication is that engineering education should clarify what forms of LLM assistance remain consistent with honest academic self-representation. If students use LLMs in drafting or revising work, instructors should make clear which uses are permissible and which are not. They may also require students to disclose how the LLM was used and what parts

of the work remain their own. This would help align LLM use with truthfulness by making students accountable for how their work represents their understanding and contribution.

A third implication is that engineering education should explicitly teach students to critically evaluate the outputs of LLMs. As previous work has shown, LLMs can produce confident but incorrect solutions in engineering contexts, and engineering students do not always reliably detect these hallucinations [4]. Therefore, students must learn to assess the validity of outputs. Engineering education should emphasize that the use of LLMs does not reduce the need for independent deliberation and judgment.

A fourth implication is that course design should discourage habitual dependence on LLMs. One way to do this is to structure assignments so that students first attempt each problem independently, then can optionally use the LLM in a limited way on a later attempt. For example, if a homework system allows multiple attempts for each problem, students could use their first attempt to work through the problem on their own before turning to an LLM for assistance on a second or third attempt. This would allow students to engage with the difficulty of the problem before seeking external support, without being penalized for making an initial independent effort.

A fifth implication is that engineering education should also identify ways to use LLMs that positively support learning. When used in a limited and reflective way, LLMs can help students clarify concepts and obtain feedback while still remaining responsible for deliberation and judgment. Such use can support practical wisdom by helping students think through problems, support truthfulness when the student's contribution remains clear, and support temperance when the LLM is used with restraint. For example, students might use an LLM to explain why a particular solution method is appropriate or to generate additional practice problems. Engineering education should therefore not treat LLMs only as risks to be controlled, but also as tools that can support learning.

VIII. CONCLUSION

This paper examines how student use of LLMs influences the development of practical wisdom, truthfulness, and temperance in engineering education. When LLMs support student deliberation and judgment without replacing them, they may complement the development of practical wisdom. When students represent their work honestly, LLM use need not undermine truthfulness. When LLMs are used with moderation rather than dependence, they may be compatible with temperance. However, when LLMs replace deliberation and judgment, distort self-representation, or encourage excessive reliance, they may reduce the intellectual and moral development that engineering education seeks to cultivate. The ethical evaluation of LLM use should therefore focus not only on whether LLMs are present, but on how they shape virtue.

REFERENCES

- [1] E. Kasneci *et al.*, "ChatGPT for good? On opportunities and challenges of large language models for education," *Learning and Individual Differences*, vol. 103, 2023.
- [2] E. R. Mollick and L. Mollick, "The impact of generative AI on education: Evidence from ChatGPT," *SSRN Electronic Journal*, 2023.
- [3] ABET, *Criteria for Accrediting Engineering Programs*. Baltimore, MD, USA: ABET, 2023.
- [4] Z. Deutsch and S. Oca, "Evaluating engineering students' detection of hallucinations in large language model solutions," presented at the ASEE Annual Conference & Exposition, 2026. [Online]. Available: https://zacharydeutsch.com/ASEE_paper.pdf
- [5] Aristotle, *Nicomachean Ethics*. T. Irwin, Trans., 2nd ed. Indianapolis, IN, USA: Hackett Publishing, 1999.
- [6] R. Bommasani *et al.*, "On the opportunities and risks of foundation models," Stanford CRFM, 2021.
- [7] T. B. Brown *et al.*, "Language models are few-shot learners," in *Advances in Neural Information Processing Systems*, 2020.
- [8] W. Zhao *et al.*, "A survey of large language models," arXiv:2303.18223, 2023.
- [9] E. M. Bender *et al.*, "On the dangers of stochastic parrots: Can language models be too big?" in *Proceedings of FAccT*, 2021.
- [10] J. R. Searle, "Minds, brains, and programs," *Behavioral and Brain Sciences*, vol. 3, no. 3, pp. 417–424, 1980.
- [11] D. C. Dennett, *The Intentional Stance*. Cambridge, MA, USA: MIT Press, 1987.
- [12] H. Petroski, *To Engineer Is Human: The Role of Failure in Successful Design*. New York, NY, USA: Vintage Books, 1992.
- [13] C. Mitcham, *Thinking Through Technology: The Path Between Engineering and Philosophy*. Chicago, IL, USA: University of Chicago Press, 1994.
- [14] J. Dunne, *Back to the Rough Ground: Practical Judgment and the Lure of Technique*. Notre Dame, IN, USA: University of Notre Dame Press, 1993.
- [15] B. S. Bloom, M. D. Engelhart, E. J. Furst, W. H. Hill, and D. R. Krathwohl, *Taxonomy of Educational Objectives: The Classification of Educational Goals*. New York, NY, USA: Longmans, Green, 1956.