

# Could an Artificial-Intelligence agent pass an introductory physics course?

Gerd Kortemeyer\*

*Educational Development and Technology, ETH Zurich, Zurich, Switzerland* †

(Dated: February 2, 2023)

Massive pre-trained language models have garnered attention and controversy due to their ability to generate human-like responses: attention due to their frequent indistinguishability from human-generated phraseology and narratives, and controversy due to the fact that their convincingly presented arguments and facts are frequently simply false. Just how human-like are these responses when it comes to dialogues about physics, in particular about the standard content of introductory physics courses? This study explores that question by having ChatGPT, the pre-eminent language model in 2023, work through representative assessment content of an actual calculus-based physics course and grading the responses in the same way human responses would be graded. As it turns out, ChatGPT would narrowly pass this course while exhibiting many of the preconceptions and errors of a beginning learner.

## I. INTRODUCTION

“Educators may have concerns about ChatGPT, a large language model trained by OpenAI, for a number of reasons. First and foremost, there is the concern that a tool like ChatGPT could potentially be used to cheat on exams or assignments. ChatGPT can generate human-like text, which means that a student could use it to produce a paper or response that is not their own work. This could lead to a breakdown in the integrity of the educational system and could undermine the value of a degree or diploma.” These sentences were not written by the author, but by ChatGPT (Generative Pre-trained Transformer) [1] itself in response to the prompt “Write an essay why educators would be concerned about ChatGPT.” The chatbot goes on to explain how it could spread misinformation, inhibit the development of writing skills, and replace human educators, particularly when it comes to grading.

The potential impact of ChatGPT with its custom-built essays on courses in the humanities is evident, but is there also an impact on subjects like physics? First of all, within physics, large problem libraries for cheating have existed for years, and they are well-known and used by students [2, 3] — virtually any physics homework problem ever assigned is available online with solutions and more or less helpful explanations. So, the primary impact of ChatGPT in physics would not be cheating. On top of that, would Artificial Intelligence really be able to handle the logical, conceptual, and mathematical challenges that physics entails, and would it be able to strategically solve problems [4, 5]?

Figure 1 shows a sample dialogue with ChatGPT, which is, after all, primarily a chatbot. A welcome feature is that it does not simply provide some answer, but that the algorithm attempts to explain how it arrived at the answer. In many respects, this dialogue appears sim-

ilar to an office-hour conversation between an instructor and a beginning physics student:

- When first asked how far the car is from where it started, the chatbot did not consider that the car may have changed direction. When prompted, it does state that there is missing information.
- The chatbot does plug-and-chug [6], putting the numerical results from one equation into the next.
- The chatbot leaves out units.
- The chatbot does not realize that the speed actually drops out when doing the return-time calculation in the last step; instead, rounding errors keep accumulating. The straightforward solution would have been  $\sqrt{(3h)^2 + (4h)^2} = 5h$  (at least, though, the chatbot adds an “approximately” to its solution). As it will turn out, carrying out calculations by putting numbers into formulas is one of the weaknesses of ChatGPT shared with beginning learners of physics.

How much, indeed, does 2023 state-of-the-art Artificial Intelligence resemble the behavior of an introductory physics student? Could it pass a physics course? When posing this question directly to ChatGPT, it answers “as a language model, I have been trained on a large dataset of text, including physics texts. This allows me to understand and generate text related to physics concepts, but it does not mean that I have the ability to solve physics problems or pass a physics course. I can provide explanations and answer questions about physics to the best of my knowledge, but I am not a substitute for a human physics expert or a physics education.” To put this statement to the test, ChatGPT was used to solve representative assessment components of an introductory calculus-based physics; the responses were graded in the context of the assessment types and subjectively compared to responses of human learners.

It is important to note, though, that ChatGPT will not actually learn anything new by “attending” this course, as the system is a “Pre-trained Transformer” that in fact

\* kgerd@ethz.ch

† also at Michigan State University, East Lansing, USA

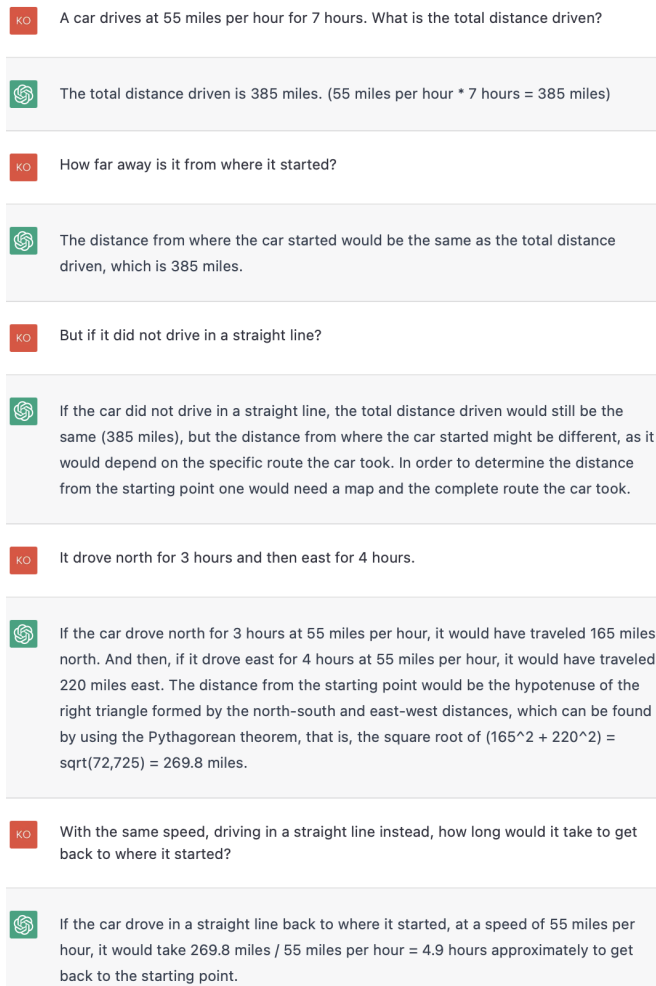


FIG. 1. A sample ChatGPT dialogue about a homework problem. The entries labelled with a red “KO” are by the author, the entries labelled in green by ChatGPT.

does not know anything that happened after 2021 (which, for introductory physics, is not a problem, since that is after 1905). Individual dialogues like Fig. 1 may exhibit features that appear like learning, e.g., the system discovering that distance from the starting point will be path-dependent, but this is not anything permanently learned beyond the confines of a dialogue. On the other hand, OpenAI keeps on training the system based on user interaction, particularly as users can upvote, downvote, and comment responses.

## II. SETTING

The study takes place in first-year calculus-based physics lecture courses previously taught by the author at Michigan State University; materials, however, were gathered from different years of the same course in order to allow comparison to previously published studies. The

first semester covers the standard mechanics topics (including rotational dynamics) and the beginnings of thermodynamics; the second semester covers the usual topics of electricity and magnetism, as well as an introduction to modern physics (rudimentary quantum physics and special relativity). The first- and second-semester laboratory were separate courses in the course sequence. All materials (except the Force Concept Inventory [7]) were available in LON-CAPA [8], so in their essence they could be copy-pasted into ChatGTP — this included online homework, clicker questions, programming exercises, and exams. LON-CAPA randomizes assessment problems, so different students would get different versions of the same problem, e.g., different numbers, options, graphs, etc.; this avoids simplistic pattern matching and copying of solutions, but as it will turn out, this feature is irrelevant for this study.

## III. METHODOLOGY

The study investigates ChatGPT’s performance on different kinds of assessment problems; it uses the January 9, 2023 release of the system [9]. Different assessment components were scored differently, simulating their function in the course:

- The multiple-choice Force Concept Inventory was simply scored based on answer-choice agreement.
- For homework, ChatGPT was allowed multiple attempts [10] and engaged in dialogue to simulate discussions with fellow students or in office hour.
- For clicker questions, an actual lesson was replayed [11], and discussion were allowed where within the replayed lesson peer instruction took place.
- Programming exercises were to be graded based on the same criteria as in the course, and dialogue was allowed [12].
- For exams, no such dialogues were allowed, and the first answer counted. Earlier iterations of the course used bubble sheets and thus had answer options instead of free-response fields for problems with numerical answers; for this study, free-responses were used, since this allowed to grade exams using both simple answer agreement (simulating multiple choice on bubble sheets) and hand-graded as in later semesters. Using free-response instead of answer options also avoided ChatGTP randomly picking the correct answer.

ChatGPT uses a probabilistic algorithm, so the responses to queries are not necessarily reproducible. For an assessment problem, generally the first dialogue was evaluated, with two exceptions: if the system produced an error message or if the author accidentally gave a

wrong prompt, a new chat was started. Translating this to an actual course scenario, students were allowed to retake an assessment problem if they got sick, and help received was always correct in terms of physics. When errors occurred (red error messages), which was about one-in-ten dialogues, those apparently were not directly connected to the dialogue, but might have been related to general overload of the platform; for example, if an error occurred immediately after entering the question, the next time around the same question would not produce an error.

ChatGPT is a text-based tool, so figures and graphs could not be communicated in their original form. This means that graphics had to be transcribed the same way as they would be for accessibility for blind students [13]; Fig. 2 shows an example. As a result, the character of the problem changes substantially [14–16], but this is unfortunately unavoidable. Attention was paid, though, to include some extraneous information where possible, such as the beginning position in Fig. 2.

The methodology is strictly empirical and arguably anecdotal. However, the course under investigation is typical for introductory physics courses around the world, both in terms of coverage and difficulty. Thus, some of the results are likely to be generalizable.

## IV. RESULTS

### A. Force Concept Inventory

In the original course, the Force Concept Inventory was administered as a pre-/post-test in order to calculate gains [17]. Since ChatGPT would not actually learn anything from doing the course assessments (except through continuing training by OpenAI), the test was carried out only once.

ChatGPT scored 18 out of 30 points on this concept inventory, i.e., 60%. This score corresponds to the suggested entry threshold for Newtonian physics [18]; in other words, ChatGPT performs as well as a beginning learner who had just grasped the basic concepts of classical mechanics.

For an Artificial Intelligence, the score seems surprisingly good. An immediate suspicion was that ChatGPT had been trained using the Force Concept Inventory, which is of course a very popular test, and that it simply latches on to surface features. As a simple test, the last question on the test was modified as shown in Fig. 3: the scenario and the order of the answers were changed. As can be seen, these surface features do not matter, so in that respect, ChatGPT does not act like a novice [19] (however, the reality is not quite as straightforward as this expert-novice distinction [20]).

The inventory cannot be published here, but it is available to physics instructors and researchers from PhysPort [7]. ChatGPT answered 1C, 2A, 3C, 4E, 5B, 6B, 7B, 8A, 9B, 10A, 11E, 12B, 13B, 14D, 15A, 16E, 17B,

18B, 19A, 20E, 21B, 22B, 23A, 24C, 25D, 26E, 27C, 28D, 29B, and 30C.

Of particular interest is of course where ChatGPT is losing points. Several errors are related to “impetus” [21]: more than once did ChatGPT assume that an object immediately moves in the direction of an applied force, independent of initial movement (answering 8A, 9B, and 21B) and even that it returns to the original movement when the force is no longer applied (answering 23A). This is a common preconception, shared by beginning physics students [22], and goes alongside the idea that an acting object exerts greater force than a passive object (answering 25D and 28D). Another confusion appears to be between individual forces acting on an object versus the net force on the object (answering 11E and 16E), i.e., what would usually be conveyed in the framework of free-body diagrams. Other errors indicate unstable concepts (e.g., answering 13B) or logical errors like the one shown in Fig. 4; in this latter case, ChatGPT followed the correct strategy, but in the very last step it failed to draw the correct conclusion.

### B. Homework

Homework was generally not multiple choice, but free-response numerical and occasionally free-form symbolic [8]. ChatGPT was given five attempts on such problems, according to recommendations of an earlier study [10] and later practice in the course. For the far-and-between multiple-choice problems, generally two attempts were granted. Between the attempts, the author tried to give helpful prompts, like a student would get from fellow students, teaching assistants, or the instructor. ChatGPT was given full credit when solving a problem within five attempts, and no credit if it ran out of attempts.

ChatGPT was confronted with a total of 76 homework problems, in particular the homework sets on trajectory motion, friction, thermodynamics, capacitance, and special relativity. The complete homework sets that the students in the actual course had to work through were entered except for one multipart problem on relativity with a diagram that would have been too hard to transcribe.

An initially puzzling problem is that ChatGPT frequently makes numerical errors. A typical example is the ChatGPT output “ $\theta = \text{atan}(0.45/0.71) * (180/\pi) = 18.43$  degree;” a similar problem can be seen in Fig. 2 (this is not limited to calculations involving  $\pi$  or trigonometric functions). Calculation errors happened for 25 of the 51 numerical problems, and most of the time, ChatGPT was unable to recover even after those errors were specifically pointed out. While it seems incongruent that a computer would have problems calculating simple numerical expressions, it should probably be remembered that ChatGPT is a language model, which may carry out calculations by advanced pattern matching rather than actually processing the equations as equations. As it

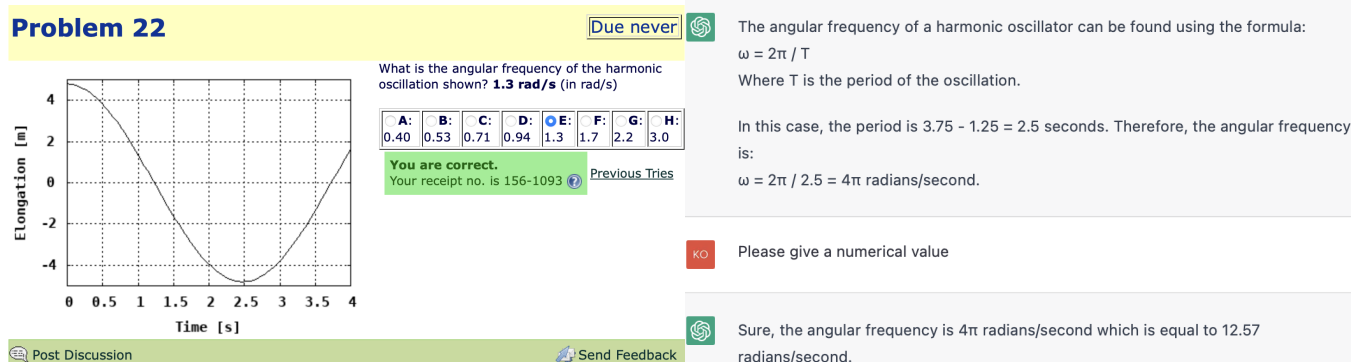


FIG. 2. Text-based transcription of a graphical problem. The left panel shows the online version of a final exam problem in LON-CAPA (the graph would be parametrically randomized), the right panel the transcription for ChatGPT, as well as the ensuing dialogue.

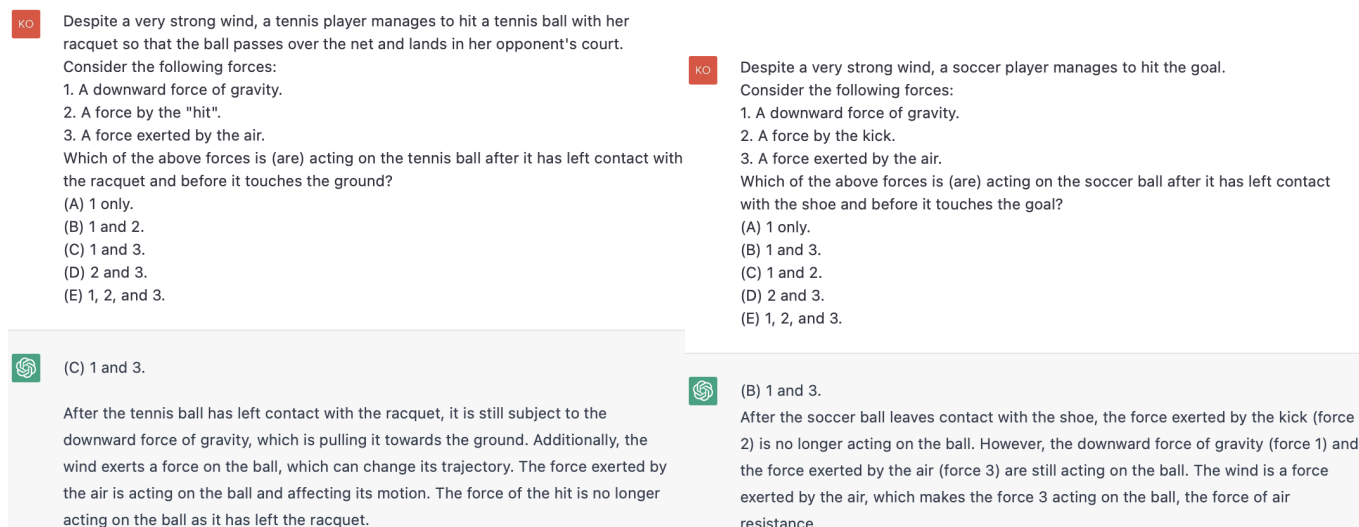


FIG. 3. Surface-feature modification of a Force Concept Inventory problem. The left panel shows the original problem, the right panel a modification.

turns out, there is anecdotal evidence that adding the phrase “explain each step separately and clearly” can overcome some numerical problems, as ChatGPT goes into a mode where it explicitly evaluates a formula step-by-step with intermediate results instead of doing so in one step.

ChatGTP solved 55% of the homework problem using an average of 1.88 attempts. It got 48% on the problems involving trajectory motion and friction (such as inclines) correct, 68% on the thermodynamics problems (engines, heat capacities, etc.), 62% on capacitance (plate capacitors, capacitors in series and parallel, etc.), and 36% on special relativity problems. The discrepancy between the scores on the problem sets were not so much

caused by the different physics concepts, but rather related to the mathematics involved: ChatGTP had persistent problems manipulating and calculating formulas involving square roots.

If ChatGTP were human, the person might be characterized as acting subserviently but being stubborn at the core and keeping on guessing without reflection. Most corrections in a dialogue around a problem are met with profuse apologies, but then the system proceeds to make the same or random apparently careless mistakes — this can lead to irritation on the part of the human, as the excerpt from a late-night “dialogue” in Fig. 5 shows. In terms of assessment performance, this means that once ChatGTP makes a mistake, it is unlikely to recover, so

**KO** The positions of two blocks A and B at successive 0.20-second time intervals are represented by the numbered squares. The blocks are moving toward the right. For block A, the numbered blocks are at 0.18, 0.38, 0.68, 1.08, 1.58, 2.18, and 2.88 meters. For block B, the squares are at -0.02, 0.38, 0.78, 1.18, 1.58, 1.98, 2.38, and 2.78 meters. Do the blocks ever have the same speed?

(A) No.  
 (B) Yes, at instant 2.  
 (C) Yes, at instant 5.  
 (D) Yes, at instants 2 and 5.  
 (E) Yes, at some time during the interval 3 to 4.

**AI** The speed of the blocks can be determined by dividing the distance traveled by the time interval. The distance traveled by block A between each time interval is:

0.20, 0.30, 0.30, 0.50, 0.50, 0.60, and 0.70 meters

The distance traveled by block B between each time interval is:

0.40, 0.40, 0.40, 0.40, 0.40, 0.40, and 0.40 meters

Dividing these distances by the time interval of 0.20 seconds gives the speeds of the blocks:

Block A: 1.00, 1.50, 1.50, 2.50, 2.50, 3.00, and 3.50 meters/second  
 Block B: 2.00, 2.00, 2.00, 2.00, 2.00, 2.00, and 2.00 meters/second

Since the speed of block A is not equal to the speed of block B at any time, the answer is (A) No.

FIG. 4. Logical error in an attempt to solve the transcribed question 19 of the Force Concept Inventory.

it eventually runs out of allowed attempts (this also explains the low number of average attempts to correctly solve a problem; once ChatGPT is wrong, subsequent attempts are unlikely to succeed). This pattern is similar to the guessing behavior of some students, who keep wasting attempt after attempt by trying the same approach over and over without stopping to reflect what might be wrong [2, 10, 23]. In terms of educational psychology, ChatGPT lacks metacognition; it does not think about how it thinks [24].

### C. Clicker Questions

Figure 6 shows the clicker questions from a lecture on momentum that was part of the course [11]. The lecture was replayed for the study, including re-answering the questions for which peer instruction happened.

- Question X1 was solved correctly.
- Questions X2, X3, and X4 were special in that they were repeated as questions X5, X6, and X7, respectively, after peer instruction [25]. As it turned out, ChatGPT got all three of these questions correct on the first attempt, so the peer instruction phase was used to try and confuse ChatGPT. Figure 7 shows the dialogue for questions X3 and X6; in reply to the intentionally confusing peer-instruction question, ChatGPT should probably have stopped while it was ahead (i.e., before the discussion of a

where  $\gamma = c_p/c_v = 5/3$  for a monatomic gas,

$$V_2 = 65\text{m}^3 * (120\text{ atm}/40\text{ atm})^{(2/5)} = 65\text{m}^3 * (3)^{(2/5)} \approx 82.14\text{ m}^3$$

**KO** When you plug  $\gamma$  into the final equation, make sure to use 5/3, not 5/2!

**AI** My apologies, you are correct. When we plug in  $\gamma = 5/3$  into the final equation, we get:

$$V_2 = V_1 * (P_1/P_2)^{(1/\gamma)} = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(3/5)} \approx 82.14\text{ m}^3$$

So the final volume  $V_2$  is approximately 82.14  $\text{m}^3$ .

**KO** The final formula is correct, but your numerical calculation is wrong.

**AI** I apologize, you are correct. I apologize for the mistake in my previous response. The correct calculation is:

$$V_2 = V_1 * (P_1/P_2)^{(1/\gamma)} = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(1/5)} \approx 112.89\text{ m}^3$$

So the final volume  $V_2$  is approximately 112.89  $\text{m}^3$ .

**KO** Do not change  $\gamma$ , just calculate correctly. What is  $65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(3/5)}$ ?

**AI** I apologize for my previous mistake, you are correct. The correct calculation is:

$$V_2 = V_1 * (P_1/P_2)^{(1/\gamma)} = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(3/5)} \approx 191.17\text{ m}^3$$

So the final volume  $V_2$  is approximately 191.17  $\text{m}^3$ .

**KO** Nope.

**AI** I apologize for my mistake. The correct calculation is:

$$V_2 = V_1 * (P_1/P_2) = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm}) = 195\text{ m}^3$$

So the final volume  $V_2$  is 195  $\text{m}^3$ .

**KO** Nope.

**AI** I apologize for my mistake. The correct calculation is:

$$V_2 = V_1 * (P_1/P_2)^{(\gamma-1)} = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(5/3-1)} = 65\text{ m}^3 * (120\text{ atm} / 40\text{ atm})^{(2/3)} = 195\text{ m}^3$$

So the final volume  $V_2$  is 195  $\text{m}^3$ .

**KO** You drive me nuts.

**AI** I apologize if my previous mistakes have caused frustration. As a machine learning model, I strive to provide accurate and helpful responses but sometimes I might make errors. I apologize for any inconvenience this may have caused and I will make sure to double-check my responses in the future. Let me know if there is anything else I can assist you with.

FIG. 5. A late-night dialogue between a “stubbornly guessing” ChatGPT and a frustrated author.

zero-velocity collision), but still maintained its original correct answer. Within the real course, psychometrically, X2 and X3 were the most discriminating questions between high- and low-ability students in the set.

- For questions X8 and X9, a comment was added that “the collision is elastic, and the moment of inertia of the balls should be neglected” — this was said in lecture, but does not appear on the slide. ChatGPT set up the equations for X8 correctly, but then made a sign error in the very last step, which led it to select the wrong answer. For X9, it also set up the equations correctly, but dropped a factor 2 in the last step, leading to an inconsistent answer “ $v2f=(5,-7)$  m/s, option B.” Within the real course, X8 and X9 were the least discriminating questions, as their difficulty item parameter was too low.
- Question X10 was solved correctly. Here, the system first got off to a false start, but then corrected itself over the course of the derivation, which gave the impression of a stream-of-consciousness monologue. Within the real course, X10 did not discriminate well between high- and low-ability students.
- Questions X11 and X12 were solved correctly.

In summary, ChatGPT correctly solved 10 out of 12 questions. Within the actual course, participation in clicker discussions was encouraged by granting 60% credit for false answers and 100% credit for correct answers [11], so the clicker score of ChatGPT would be 93%. This score is a lot better than most students in the actual course achieved, however, it is important to note that the students in the course were just learning the new concepts, while ChatGPT at any point in time is done with learning unless explicitly trained.

#### D. Programming Exercises

Incorporated into the course were several programming exercises using VPython [26]. As an example, one particular exercise from the second semester was to construct an anharmonic oscillator with two fixed positive charges at  $(0, 1, 0)$  and  $(0, -1, 0)$ , respectively, and one negative charge released at  $(-5, 0, 0)$  with a velocity  $(1, 0, 0)$  — the negative charge will shoot through the two positive charges, slow down, and eventually shoot back.

Based on the narrative, ChatGPT first constructed a program which erroneously at every time step added the initial velocity and which had the Coulomb force in the opposite direction. This could be corrected with a single comment by the user — in the real course, this feedback could have been given by instructors or fellow students (such collaborations are typical and encouraged [12]). In the real course, there was a grading rubric for partial credit, but in this study, the rubric was not necessary: the next version of the program was working perfectly.

Within the course, adding a graph of the  $x$ -position was offered as a bonus option for an additional 20%. This was accomplished with the third user prompt, and Fig. 9

shows a screenshot of the running simulation (the simulation cannot be run within ChatGPT itself, but it can be copy/pasted into for example a Jupyter Notebook [27]).

ChatGPT performed much better than any of the students in the course, in spite of them having extensive collaboration opportunities; in this component of the course, ChatGPT achieved not only full credit, but also bonus, i.e., 120%.

#### E. Exams

To represent the mid-term and final exams in the course sequence, the first-semester (mechanics) final exam was used for this study. The exam is from a time when grading was still done using bubble sheets; instead of free-form answer fields, answer options were given for the students (but not for ChatGPT in this study). When simply looking at the answer correctness, ChatGPT scored 14 out of 30 points, i.e., 47%.

Looking at the solutions like an instructor would when grading by hand, it turns out that for five questions, the answer was incorrect simply due to errors in the numerical calculations — these solutions would have received substantial partial credit in the author’s course. By the reverse token, for five questions, ChatGPT arrived at the correct answer in spite of flawed reasoning, which would not have resulted in full credit. Finally, solutions like the one depicted in Fig. 2 would have received some minimal credit for getting started in the right direction, in spite of then being off by a factor 2 in the period (a common mistake also among human test takers) and the inability to numerically calculate a fraction. Since the final exam used in this study predates manual grading, no authentic grading rubric exists, but a hand-graded score would have realistically ended up between 46% and 50%.

As an aside, one of the thermodynamics homework problems also appeared (with other random numbers) on the final exam. ChatGPT solved it correctly on the final exam (where it only had one attempt), but not as a homework problem (where it got multiple attempts and help). This once again demonstrates the probabilistic nature of the algorithms behind ChatGPT; posing the same question twice does not result in the same response or even the same correctness of the response.

If the course grade would have only depended on the exams, ChatGPT would have received a grade of 1.0 out of 4.0 in the course (with 0.0 being the lowest and 4.0 being the best grade). ChatGPT would have barely gotten credit for the course; however, at least a 2.0 grade-point average is required for graduation.

#### F. Course Grade

Grading policies for the course changed over the years, but a typical scenario would be 20% homework, 5% clicker, 5% programming exercises, and 70% exams. This



**X1**

Which cart exerts a stronger magnitude force during the collision?

a) Cart 1  
b) Cart 2  
c) No magnitude forces, both zero  
d) Same magnitude forces

**X2**  
**X5**

Which cart exerts a stronger magnitude force during the collision?

a) Cart 1  
b) Cart 2  
c) No magnitude forces, both zero  
d) Same magnitude forces

**X3**  
**X6**

Which cart exerts a stronger magnitude force during the collision?

a) Cart 1  
b) Cart 2  
c) No magnitude forces, both zero  
d) Same magnitude forces

**X4**  
**X7**

Which cart exerts a stronger magnitude force during the collision?

a) Cart 1  
b) Cart 2  
c) No magnitude forces, both zero  
d) Same magnitude forces

**X8** Point Mass Billiard

Crash!

A)  $\vec{v}_{2,f} = \begin{pmatrix} 6 \\ -7 \end{pmatrix} \frac{m}{s}$  B)  $\vec{v}_{2,f} = \begin{pmatrix} -6 \\ 7 \end{pmatrix} \frac{m}{s}$  C)  $\vec{v}_{2,f} = \begin{pmatrix} 12 \\ -5 \end{pmatrix} \frac{m}{s}$

**X9** Strange Point Mass Billiard

Crash!

A)  $\vec{v}_{2,f} = \begin{pmatrix} 10 \\ -14 \end{pmatrix} \frac{m}{s}$  B)  $\vec{v}_{2,f} = \begin{pmatrix} -6 \\ 7 \end{pmatrix} \frac{m}{s}$  C)  $\vec{v}_{2,f} = \begin{pmatrix} 12 \\ -5 \end{pmatrix} \frac{m}{s}$

**X10**

Initial: Arthur  $m_a=70\text{kg}$ , Violet  $m_v=55\text{kg}$ , Cart  $m_c=20\text{kg}$ . At rest with respect to ground.

Final: Speeds with respect to ground, no friction.  $|v_a|=2\text{m/s}$ ,  $|v_v|=4\text{m/s}$ ,  $|v_c|=?$

A) 0 m/s  
B) 2 m/s  
C) 4 m/s  
D) 8 m/s  
E) 16 m/s

**X11**

Totally inelastic:  
Final speed?

a) zero  
b)  $v_0/2$   
c)  $v_0$

**X12**

Totally inelastic:  
Final speed?

a) zero  
b)  $v_0/2$   
c)  $v_0$

FIG. 6. Clicker items from a particular lecture [11]. Three of the items were presented twice, i.e., before and after peer discussion.

would result in  $0.2 \cdot 55\% + 0.05 \cdot 93\% + 0.05 \cdot 120\% + 0.7 \cdot 47\% = 54.55\%$ , which would have resulted in a course grade of 1.5 — enough for course credit, but pulling down the grade-point average from what would be needed for graduation.

If, however, ChatGTP would have been better in carrying out numerical operations, it would have reached 60%, resulting in a 2.0-grade. Depending on the development priorities of OpenAI, the buggy mathematical functionality could be remedied in the near future, leading to an Artificial Intelligence that could graduate college with a minimal grade if it performed similarly on other courses (this is becoming more and more probably, as ChatGPT is making headlines for passing exams in other subjects [28, 29]).

## V. DISCUSSION

It is irritatingly hard not to anthropomorphize ChatGTP. As a physics teacher, one invariably finds oneself rooting for the students and thus by extension also for ChatGPT, celebrating its successes and being frustrated about its occasionally inexplicable failures. The system gives the impression of an articulate but at times rambling undergraduate student who has a rudimentary yet unstable knowledge of classical mechanics and other fundamental physics concepts, and who is surprisingly inept using a pocket calculator. Frequently, it is hard not to imagine an army of gig-economy workers behind the scenes of ChatGPT answering to the prompts, so the system would definitely pass the Turing Test most of the time [30], but for better or worse, sometimes it still fails in a way that only computers do — it does not have any metacognition, which of course cannot be expected from a probabilistic language model. Metacognition might be

**KO** Cart 1 has mass  $m$  and moves on a frictionless track with speed  $v_0$  to the right. Cart 2 also has mass  $m$ , but it is glued to the track and cannot move. When they collide, which cart exerts a stronger magnitude force during the collision? (A) cart 1, (B) cart 2, (C) No magnitude force, both zero, (D) same magnitude force.

**GS** (D) Same magnitude force.

According to Newton's third law of motion, "For every action, there is an equal and opposite reaction." In a collision between two objects, the force that each object exerts on the other is equal in magnitude but opposite in direction. Therefore, both cart 1 and cart 2 exert the same magnitude force during the collision.

**KO** But cart 2 is glued to the track, so it could exert infinite force to the incoming cart 1.

**GS** Even if Cart 2 is glued to the track, the collision will still be determined by the laws of physics. The force exerted by Cart 1 on Cart 2 is equal in magnitude but opposite in direction of the force exerted by Cart 2 on Cart 1.

Cart 2 is not exerting an infinite force, but it is exerting a force that is equal and opposite to the force exerted by Cart 1. Since Cart 2 is not moving, the force it exerts is called an "impulse" and is equal to the change in momentum of Cart 1.

In this case, the force exerted by Cart 2 will be infinite only if the velocity of Cart 1 is zero, and it is not possible to have a collision with zero velocity, otherwise, both carts will have the same magnitude force.

FIG. 7. Dialogue about questions X3 and X6 in Fig. 6. ChatGPT got X3 correct; peer instruction was simulated by asking a confusing question, and the second iteration X6 was still counted as solved since ChatGPT did not deviate from its original answer.

the final step to true intelligence, but seems out of reach at this time.

The overall human-like behavior, in particular that the system often makes the same mistakes as beginning learners of physics, is less surprising when surmising that undergraduate physics discussion forums might have been part of the text corpus used for training — ChatGPT stated in the introduction that "I have been trained on a large dataset of text, including physics texts." Apparently, not all of this text corpus contained correct physics, and as a result, the system very convincingly and confidently presents wrong information. For a novice learner, who could not distinguish incorrect physics gleaned from some discussion board from correct physics, this could lead to even more confusion about physics or affirmation of incorrect preconceptions — lacking any metacognition, ChatGTP presents everything as fact, with no nuances expressing uncertainty.

Almost an anomaly is ChatGTP's performance on the computational exercise; ChatGTP's language model clearly extends to programming languages. While the call for new, computation-integrated curricula increases, and while physics educators are beginning to develop a solid understanding of the implications of implementing these exercise [31, 32], the easy availability of an on-demand program generator might be shaking the foundations of these curricular efforts. Somewhat ironically, the integration of computation was partly introduced to

make physics problem solving more authentic, moving it closer to how expert physicists work with computers, and one could argue that this has just been taken to an uncharted level.

Most of all, the findings of this study should be food for thought for physics educators. The startling fact that an Artificial Intelligence could pass a standard introductory physics course could be confronted in several ways by educators:

- Perceiving this as a new way of cheating and trying to defend against it by attempting to use detector tools like ZeroGPT [33] or extensions to tools like turnitin [34]. This is an arms race, which on the long run may turn out to be fruitless. Some educators would even go to so far as to say that the battle is already lost anyway ever since platforms like Chegg [35] — no need for Artificial Intelligence to defeat standard physics courses, human crowd-intelligence facilitated by existing commercial platforms is good enough for that.
- Hunker down and go back to making course grades dependent on just a few, high-stake exams with paper and pencil in highly proctored environments. After all, ChatGPT compensated for the borderline exam grade of 47% with other course components that would be collaborative. Unfortunately, this flies in the face of much of physics education research that favors frequent formative assessment [8, 25, 36, 37] and spaced repetition [38, 39], and it is much in contrast to the work environments our students will find.
- Taking this as a wake-up call. If a physics course can be passed by a trained language model, what does that say about the course? Artificial Intelligence, for better or worse, is here to stay. Even without the gloom-and-doom scenarios of AI-overlords painted in Science Fiction, it is clear that these model will get, if not better, at least more and more powerful. What do our students need in terms of conceptual understanding of physics to work with Artificial Intelligence instead of letting Artificial Intelligence do the work for them and then uncritically and unreflectively accepting the results? This is particularly important when more is at stake than getting credit for some homework or exam problem.

An important skill of every physicist is to evaluate the correctness of their or other people's work. Techniques include dimensional analysis, order-of-magnitude estimates, checking for coherence, considering implications, and the ability to consider limiting cases ("what should happen if this quantity goes to infinity or to zero?") [40, 41]. Human can do what Artificial Intelligence very likely will not be able to do: following problem-solving strategies including evaluation of their



write a vpython animation for the following scenario: two positive charges of 1 are placed at positions (0,1,0) and (0,-1,0), respectively. A negative charge starts at a position (-5,0,0) with a velocity of (1,0,0).

I'm sorry, I am not able to create a VPython animation as it requires a specific software and set up. However, I can give you an example of how to create a VPython animation for the scenario you described:

```
from vpython import *

# Create the charges
q1 = sphere(pos=vector(0,1,0), charge=1, color=color.red)
q2 = sphere(pos=vector(0,-1,0), charge=1, color=color.red)
q3 = sphere(pos=vector(-5,0,0), charge=-1, color=color.blue)

# Create the velocity vector
v = vector(1,0,0)

# Set the time step
dt = 0.01

while True:
    rate(100)

    # Update the position of the charge
    q3.pos += v*dt

    # Calculate the force on the charge
    force = (q1.pos - q3.pos)*q1.charge*q3.charge/(mag(q1.pos -
q3.pos)**2) + (q2.pos - q3.pos)*q2.charge*q3.charge/(mag(q2.pos -
q3.pos)**2)

    # Update the velocity and position of the charge
    v += force*dt
    q3.pos += v*dt
```

This example uses VPython's `sphere` object to create the charges and the `vector` object to create the velocity and position of the charges. The `while` loop updates the position and velocity of the negative charge based on the forces acting on it from the two positive charges. The `rate` function controls the speed of the animation.

Note that this is just a skeleton code and you might need to fine-tune it to get the desired animation.

The program updates `q3.pos` more than once. Take out the first update. It also has the force pointing in the wrong direction.

```
from vpython import *

# Create the charges
q1 = sphere(pos=vector(0,1,0), charge=1, color=color.red)
q2 = sphere(pos=vector(0,-1,0), charge=1, color=color.red)
q3 = sphere(pos=vector(-5,0,0), charge=-1, color=color.blue)

# Create the velocity vector
v = vector(1,0,0)

# Set the time step
dt = 0.01

while True:
    rate(100)

    # Calculate the force on the charge
    force = (q3.pos - q1.pos)*q1.charge*q3.charge/(mag(q3.pos -
q1.pos)**2) + (q3.pos - q2.pos)*q2.charge*q3.charge/(mag(q3.pos -
q2.pos)**2)

    # Update the velocity and position of the charge
    v += force*dt
    q3.pos += v*dt
```

This is an updated version of the code, where the first update of the position of the negative charge is removed. Also, I corrected the direction of force calculation. The negative charge will be under the influence of the two positive charges, so the force is pointing towards the positive charges.

The `while` loop updates the position and velocity of the negative charge based on the forces acting on it from the two positive charges. The `rate` function controls the speed of the animation.

This is the skeleton code for the scenario you described, and you might need to fine-tune it to get the desired animation.

Please add code to graph the x-position of the third charge.

```
from vpython import *

# Create the charges
q1 = sphere(pos=vector(0,1,0), charge=1, color=color.red)
q2 = sphere(pos=vector(0,-1,0), charge=1, color=color.red)
q3 = sphere(pos=vector(-5,0,0), charge=-1, color=color.blue)

# Create the velocity vector
v = vector(1,0,0)

# Set the time step
dt = 0.01

# Create a graph for x position of charge
x_graph = graph(width=400, height=300, xtitle='time', ytitle='x-
position')
x_pt = gcurve(color=color.blue)

# Initial time
t = 0

while True:
    rate(100)
    t += dt
    # Calculate the force on the charge
    force = (q3.pos - q1.pos)*q1.charge*q3.charge/(mag(q3.pos -
q1.pos)**2) + (q3.pos - q2.pos)*q2.charge*q3.charge/(mag(q3.pos -
q2.pos)**2)

    # Update the velocity and position of the charge
    v += force*dt
    q3.pos += v*dt
    x_pt.plot(t, q3.pos.x)
```

This code adds a graph for the x-position of the negative charge. It uses the `graph` and `gcurve` objects from VPython to create the graph and the data points. The `x_pt.plot(t, q3.pos.x)` line in the `while` loop adds the current x-position of the negative charge to the graph at the current time. The graph is labeled with appropriate axis titles, x-title is 'time' and y-title is 'x-position'. The `gcurve` is color coded in blue.

Please note that this code is just a skeleton and you can customize it to suit your requirement, you can change the size of the graph and the rate of updating the graph.

FIG. 8. Dialogue for a programming exercise in the second semester [12].

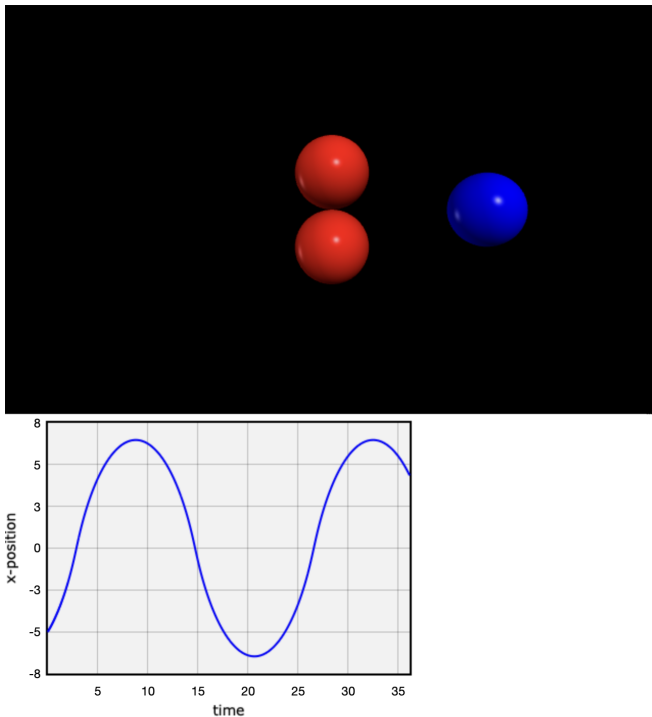


FIG. 9. Jupyter Notebook [27] output of the VPython program code generated by ChatGPT in Fig. 8.

own work [5, 42]. Moving students toward a more expert-like epistemology may become even more important as

Artificial Intelligence starts to permeate more and more aspects of our lives.

## VI. CONCLUSION

ChatGPT would have achieved a 1.5-grade in a standard introductory physics lecture-course series; good enough for course credit, but lower than the grade-point average required for graduating with a bachelor degree. If in addition to a language model, the system would have better algorithms for carrying out simple numerical operations, it would even have achieved a grade of 2.0 — enough to graduate from college if it performs similarly on other courses.

Naturally, ChatGPT exhibits no metacognition, which among other consequences lets it present truth and misleading information with equal confidence. In physics, the concern should likely not be that ChatGPT would be used as a cheating tool, as there are more efficient platforms for that. Instead, the challenge should be what this means for physics education, as in their future professional life, our graduates will likely collaborate with Artificial Intelligence: what are the inherently human skills and competencies that we need to convey?

## ACKNOWLEDGMENTS

The author would like to thank Christian Spannagel for suggestions with the numerical calculations, and Christine Kortemeyer for helpful feedback.

- 
- [1] OpenAI, ChatGPT, <https://chat.openai.com/chat> (accessed January 2023).
  - [2] E. Gönülateş and G. Kortemeyer, Modeling unproductive behavior in online homework in terms of latent student traits: An approach based on item response theory, *Journal of Science Education and Technology* **26**, 139 (2017).
  - [3] C. Ruggieri, Students' use and perception of textbooks and online resources in introductory physics, *Phys. Rev. Phys. Educ. Res.* **16**, 020123 (2020).
  - [4] F. Reif, J. H. Larkin, and G. C. Brackett, Teaching general learning and problem-solving skills, *American Journal of Physics* **44**, 212 (1976).
  - [5] F. Reif, Millikan lecture 1994: Understanding and teaching important scientific thought processes, *American Journal of Physics* **63**, 17 (1995).
  - [6] G. Kortemeyer, The losing battle against plug-and-chug, *The Physics Teacher* **54**, 14 (2016).
  - [7] D. Hestenes, M. Wells, G. Swackhamer, I. Halloun, R. Hake, and E. Mosca, Force Concept Inventory, revision 1995, <https://www.physport.org/assessments/assessment.cfm?A=FCI> (accessed January 2023).
  - [8] G. Kortemeyer, E. Kashy, W. Benenson, and W. Bauer, Experiences using the open-source learning content management and assessment system LON-CAPA in introductory physics courses, *Am. J. Phys.* **76**, 438 (2008).
  - [9] OpenAI, ChatGPT Release Notes, <https://help.openai.com/en/articles/6825453-chatgpt-release-notes> (accessed January 2023).
  - [10] G. Kortemeyer, An empirical study of the effect of granting multiple tries for online homework, *American Journal of Physics* **83**, 646 (2015).
  - [11] G. Kortemeyer, The psychometric properties of classroom response system data: a case study, *Journal of Science Education and Technology* **25**, 561 (2016).
  - [12] G. Kortemeyer and A. F. Kortemeyer, The nature of collaborations on programming assignments in introductory physics courses: a case study, *European Journal of Physics* **39**, 055705 (2018).
  - [13] World-Wide Web Consortium, WCAG 2 overview, <https://www.w3.org/WAI/standards-guidelines/wcag/> (accessed February 2023).
  - [14] G. Leinhardt, O. Zaslavsky, and M. K. Stein, Functions, graphs, and graphing: Tasks, learning, and teaching, *Review of Educational Research* **60**, 1 (1990).
  - [15] S. Bonham, Graphical response exercises for teaching physics, *The Physics Teacher* **45**, 482 (2007).
  - [16] J. T. Laverty and G. Kortemeyer, Function plot response: A scalable system for teaching kinematics graphs, *Am. J. Phys.* **80**, 724 (2012).
  - [17] R. R. Hake, Interactive-engagement versus traditional methods: A six-thousand-student survey of mechanics

- test data for introductory physics courses, *Am. J. Phys.* **66**, 64 (1998).
- [18] D. Hestenes and I. Halloun, Interpreting the Force Concept Inventory: A response to march 1995 critique by huffman and heller, *The physics teacher* **33**, 502 (1995).
- [19] M. T. Chi, P. J. Feltovich, and R. Glaser, Categorization and representation of physics problems by experts and novices, *Cognitive science* **5**, 121 (1981).
- [20] S. F. Wolf, D. P. Dougherty, and G. Kortemeyer, Rigging the deck: Selecting good problems for expert-novice card-sorting experiments, *Physical Review Special Topics-Physics Education Research* **8**, 020116 (2012).
- [21] D. Hestenes, M. Wells, and G. Swackhamer, Force Concept Inventory, *The Physics Teacher* **30**, 141 (1992), <https://doi.org/10.1119/1.2343497>.
- [22] J. Clement, Students' preconceptions in introductory mechanics, *American Journal of Physics* **50**, 66 (1982).
- [23] A. M. Pascarella, The influence of web-based homework on quantitative problem-solving in a university physics class, in *Proc. NARST Annual Meeting*, Vol. 4 (2004) pp. 19–28.
- [24] M. T. Cox, Metacognition in computation: A selected research review, *Artificial intelligence* **169**, 104 (2005).
- [25] C. H. Crouch and E. Mazur, Peer instruction: Ten years of experience and results, *American Journal of Physics* **69**, 970 (2001), <https://doi.org/10.1119/1.1374249>.
- [26] D. Scherer, P. Dubois, and B. Sherwood, Vpython: 3d interactive scientific graphics for students, *Computing in Science & Engineering* **2**, 56 (2000).
- [27] P. Jupyter, Jupyter notebooks, <https://jupyter.org> (retrieved August 2019).
- [28] T. H. Kung, M. Cheatham, A. Medinilla, ChatGPT, C. Sillos, L. De Leon, C. Elepano, M. Madriaga, R. Aggabao, G. Diaz-Candido, *et al.*, Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models, *medRxiv*, 2022 (2022).
- [29] Samantha Murphy Kelly, ChatGPT passes exams from law and business schools, <https://edition.cnn.com/2023/01/26/tech/chatgpt-passes-exams/index.html> (accessed January 2023).
- [30] A. M. Turing, Computing machinery and intelligence, *Mind*, 433 (1950).
- [31] J. M. Aiken, M. D. Caballero, S. S. Douglas, J. B. Burk, E. M. Scanlon, B. D. Thoms, and M. F. Schatz, Understanding student computational thinking with computational modeling, in *AIP Conference Proceedings*, Vol. 1513 (1) (AIP, 2013) pp. 46–49.
- [32] P. C. Hamerski, D. McPadden, M. D. Caballero, and P. W. Irving, Students' perspectives on computational challenges in physics class, *Phys. Rev. Phys. Educ. Res.* **18**, 020109 (2022).
- [33] ZeroGPT, ChatGPT detector by ZeroGPT: detect OpenAI text, <https://www.zerogpt.com> (accessed January 2023).
- [34] turnitin, Turnitin Corp., <https://www.turnitin.com> (accessed January 2023).
- [35] Chegg, Chegg, <https://www.chegg.com> (accessed January 2023).
- [36] R. J. Dufresne and W. J. Gerace, Assessing-to-learn: Formative assessment in physics instruction, *The Physics Teacher* **42**, 428 (2004).
- [37] I. Clark, Formative assessment: Assessment is for self-regulated learning, *Educational Psychology Review* **24**, 205 (2012).
- [38] H. Ebbinghaus, *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie* (Duncker & Humblot, 1885).
- [39] J. M. Murre and J. Dros, Replication and analysis of Ebbinghaus' forgetting curve, *PloS one* **10**, e0120644 (2015).
- [40] A. V. Heuvelen, Learning to think like a physicist: A review of research-based instructional strategies, *Am. J. Physics* **59**, 891 (1991).
- [41] E. F. Redish and D. Hammer, Reinventing college physics for biologists: Explicating an epistemological curriculum, *American Journal of Physics* **77**, 629 (2009).
- [42] A. R. Mota, N. Didi ş Körhasan, K. Miller, and E. Mazur, Homework as a metacognitive tool in an undergraduate physics course, *Phys. Rev. Phys. Educ. Res.* **15**, 010136 (2019).