

## Educational data augmentation in physics education research using ChatGPT

Fabian Kieser<sup>1</sup>, Peter Wulff<sup>1</sup>, Jochen Kuhn<sup>2</sup>, and Stefan Küchemann<sup>2,\*</sup>

<sup>1</sup>*Physics and Physics Education Research, Heidelberg University of Education,  
Im Neuenheimer Feld 561, 69120 Heidelberg, Germany*

<sup>2</sup>*Chair of Physics Education, Faculty of Physics, Ludwig-Maximilians-Universität München  
(LMU Munich), Geschwister-Scholl-Platz 1, 80539 Munich, Germany*



(Received 28 July 2023; accepted 3 October 2023; published 25 October 2023)

Generative AI technologies such as large language models show novel potential to enhance educational research. For example, generative large language models were shown to be capable of solving quantitative reasoning tasks in physics and concept tests such as the Force Concept Inventory (FCI). Given the importance of such concept inventories for physics education research, and the challenges in developing them such as field testing with representative populations, this study seeks to examine to what extent a generative large language model could be utilized to generate a synthetic dataset for the FCI that exhibits content-related variability in responses. We use the recently introduced ChatGPT based on the GPT 4 generative large language model and investigate to what extent ChatGPT could solve the FCI accurately (RQ1) and could be prompted to solve the FCI as if it were a student belonging to a different cohort (RQ2). Furthermore, we study, to what extent ChatGPT could be prompted to solve the FCI as if it were a student having a different force- and mechanics-related preconception (RQ3). In alignment with other research, we found that ChatGPT could accurately solve the FCI. We furthermore found that prompting ChatGPT to respond to the inventory as if it belonged to a different cohort yielded no variance in responses, however, responding as if it had a certain preconception introduced much variance in responses that approximate real human responses on the FCI in some regards.

DOI: [10.1103/PhysRevPhysEducRes.19.020150](https://doi.org/10.1103/PhysRevPhysEducRes.19.020150)

### I. INTRODUCTION

Conducting empirical educational research is at the center of learning and instruction sciences, for instance, for evaluating the effectiveness of instructional materials, proving theoretical frameworks, or designing test instruments. The design of effective learning materials or test instruments often requires a large amount of human effort to perform the studies and several iterations within the empirical validation process for the optimization of the materials. Additionally, conducting empirical studies requires financial investments, such as providing incentives for participation.

Furthermore, in the past decade, learning systems based on artificial intelligence (AI), such as AI-based intelligent tutoring systems or automated feedback systems, have received growing attention, and several efforts have targeted the implementation of AI tools in learning environments. For a good performance, such AI tools require training and fine-tuning with a large amount of empirical

data. Although there are a growing number of pretrained AI-based models that provide good results for cases similar to the training data, it is often necessary to fine-tune them for desired performance in specific applications.

Recent advances in large language models (LLM) could have several opportunities for education. Kasneci *et al.* suggest the number of opportunities for learning and teaching, such as support in writing tasks, development of critical thinking, individual learning opportunities, semiautomated grading of students' works, and lesson planning [1]. For educational research, LLM may also help with summarizing and categorization tasks. For instance, van de Shoot *et al.* developed a tool for automated title and abstract screening for systematic literature reviews [2], which was shown to save a significant amount of time. Particularly, the release of ChatGPT (based on GPT 3.5) has stimulated several research efforts to explore opportunities for education. For instance, ChatGPT can be used to summarize texts, zero-shot text classification tasks, and programming tasks [3,4].

It was also found that ChatGPT is able to solve a variety of problems with zero-shot learning, it performs reasonably well on university exams in medicine and can support in solving higher-order problems in pathology [5]. These results offer the opportunity to test developed assessments before administering it to students and subsequently

\*.kuechemann@lmu.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

optimizing them based on the answers of ChatGPT. This is an important benefit, as it can save time and effort in the assessment development process and yield insights into the psychometric properties of the instruments. Additionally, in science education research, concept tests, such as concept inventories, play an important role in diagnosing students' conceptual understanding and informing teachers about prevailing students' difficulties.

In physics education research, the most used concept inventory (CI) is the Force Concept Inventory (FCI) developed by Hestenes *et al.* [6]. As a common concept test, it consists of several multiple-choice questions that target specific preconceptions related to Newtonian mechanics, in particular the force concept. Therefore, the answer alternatives are specifically designed to reflect these preconceptions. To use AI-based methods such as LLM to test features of concept tests, e.g., if they capture students' difficulties, it is necessary that LLM are aware of these difficulties and know how students would answer as if they experienced one or more of these difficulties. Furthermore, if an LLM would be able to answer concept questions from the perspective of a student who experiences difficulties, the created synthetic data could be used to train and fine-tune smaller AI tools for education that support learners or better inform teachers. Therefore, it would be a great benefit if LLM were able to create a synthetic dataset that closely resembles a realistic student dataset, as it would save a significant amount of time and effort in the development of assessments and concept tests.

For this purpose, we explore in this work the opportunity of creating synthetic empirical data with LLM. First, we provide an overview of relevant previous works and then specify research questions that target certain aspects of the suitability to generate a synthetic dataset in physics education research.

## II. THEORETICAL BACKGROUND

### A. Concept inventories in science subjects

Physics in particular and science subjects more generally are characterized by a structured knowledge base that is centered on key concepts [7,8]. Hence, the conceptual understanding of key concepts in scientific disciplines such as evolution or force is important and key to scientific literacy. To measure high-level conceptual understanding and evaluate pedagogical practices, science education researchers devised so-called concept inventories [9,10]. Concept inventories have been developed and widely employed to assess the understanding of key concepts such as biological evolution [11], climate change [12], or Newtonian force [13]. These concept inventories typically involve multiple-choice questions, where designing distractors that capture common preconceptions is of crucial importance. Concept inventories have helped science education researchers empirically show to what extent a basic

level of conceptual understanding is reached in school and university education. Oftentimes, it is found that a full understanding of key concepts such as in Newtonian mechanics, the force concept is typically reached only by a few students even after (traditional) instruction. Moreover, concept inventories helped researchers to glean evidence supporting the effectiveness of interactive engagement instructional formats, which yielded higher learning gains [13]. The concept inventories thus helped to substantially improve instruction, e.g., of mechanics in introductory physics, across universities [14].

### B. Development of CIs

Designing concept inventories was described by Adams and Wieman as a process consisting of consecutive phases [15], namely (i) delineation of the purpose and scope of the construct to be measured, (ii) development and evaluation of test specifications, (iii) field testing, evaluation, and selection of items, and (iv) assembly and evaluation of the test. Phase 1 includes information about expert thinking on the construct and requirements for the concept inventory to be implemented, such as no required training and ease of use. In phase 2, decisions on psychometric properties are considered, as well as the target population for which the concept inventory is designed. Phases 3 and 4 comprise the bulk of the work needed to design the concept inventory. Six iterative steps are distinguished that include (a) establishing topics of interest, (b) interviews and observations to identify student thinking, (c) create open-ended survey to probe student thinking, (d) create forced answer test, (e) carry out validation interviews, and (f) administer to classes and run statistical tests. Establishing the topics of interests typically requires a resource-intensive Delphi process [10]. Often, 15 to 20 experts need to be recruited to discuss and refine topics or constructs in concept inventories. [10,15]. Developing a concept inventory is thus an extensive and time-consuming task that involves substantial resources, where AI technologies might be useful in almost any phase and step. Once qualitative, open-ended analyses are completed, suitable forced-choice (e.g., multiple-choice) items have to be field-tested where students' responses are sampled in order to determine the underlying structure of the concept inventory in terms of psychometric properties [14,15]. In the field tests, class sizes of a "few hundred or more [students] are desirable" [15] (p. 1303). In line with generalizability theory, administering parallel tests to two equivalent populations is desired [15]. However, these requirements are difficult to meet and place restrictions on small institutions that might not be able to develop concept inventories in the first place. It might also be difficult to sample representative student populations, because of various reasons such as restricted access, ecological considerations, or conceptual problems such as defining the target population. Concept inventories would also have to be field tested to different age groups to determine potential differences [14].

Given this resource-intensive process of field testing concept inventories, it would be desirable for the PER community to enhance existing strategies to field test CI. A recently evolved strategy might be the use of generative AI technologies such as LLM that were found to excel at various language generation tasks that even involved knowledge application. LLMs have been found to be capable of generating synthetic data and solving problems without being explicitly instructed to do so; therefore, they are valuable candidates for outsourcing some of the field testing in concept inventory development—and test instrument development more generally.

### C. Large language models (LLM)

AI technologies have enabled tremendous progress in fields such as computer vision and natural language processing (NLP). The first breakthrough moment for machine learning (inductive problem solving by computers) was probably when deep neural networks could be utilized to accurately classify pictures. Even more so, since its inception [16], and especially since 2017 so-called LLM have had a profound impact on solving all sorts of language-related problems as diverse as question answering, translation, or summarizing text [17–19]. There is a vast variety of LLM architectures that accomplish specific tasks. For example, sequence-to-sequence models were used to translate an input to an output, or encoder architectures were used to classify an input into predefined categories. Most often, some form of deep neural networks underlie the architecture of LLM. A general feature of these LLM was that they are sensitive to the number of parameters of the artificial neural network, train set size, or amount of compute time, where increases in all of these features generally improve the performance of the LLM [20].

Many LLM are trained in two steps. The first step consists of semisupervised learning in which a masked language model objective is pursued. In masked language modeling context, words are masked out and have to be predicted by the algorithm, given a related sequence of words [21,22]. Early LLM was trained based on large text corpora such as the Common Crawl (i.e., dump of the Internet, 60% [23]), WebText2 (22%), or Wikipedia (3%) with the training objective to predict next words in a sequence (a sort of cloze test, [18]) and classify if a subsequent sentence actually follows a given sentence [18]. To the surprise of many experts, such a simple training objective enabled large enough models to gain capabilities to perform entirely novel tasks that were not seen in the training process [24]. Some resemblance of masked language modeling to human language processing and cognition, namely predictive coding where future words are sought to be predicted given the current words were recognized [25]. The second step consists of a supervised learning task in which the pretrained model (also called the foundational model, [26]) is trained or fine-tuned on a

specific task on a labeled dataset. In this part, the LLM predicts classes based on textual information.

Among the first LLM that utilized a transformer architecture and solved a wide range of unseen tasks was the Generative Pretrained Transformer 3 (GPT 3) by OpenAI. This model was tested on a wide range of tasks such as common-sense reasoning. It was found that it achieved substantial accuracy [24]. A novel paradigm emerged, namely prompting LLM, i.e., providing it context with the input that it uses to generate a response. One prompting strategy was to provide the LLM examples of similar task solutions that could improve accuracy [24]. Some researchers contended that LLM with their tremendous size (GPT 3 has 175B parameters) overfit the training data and are essentially “stochastic parrots” [27] that regurgitate the training data. However, analyses of data contamination (i.e., evaluation examples being present in the training data) revealed that the LLM could solve tasks that were not contained in the training dataset [24].

While word analogy problems and other fairly simple problems have been solved by LLM rather early, quantitative reasoning as in disciplines such as biology, chemistry, or physics was a more difficult task to solve. AI researchers introduced more advanced training tasks that enabled LLM to solve quantitative reasoning problems as well to some extent [28]. In particular, chain-of-thought prompting was found to be particularly effective [29]. For once, researchers used science-related content to specifically further fine-tune (i.e., adapt the weights) the LLM in a learning paradigm called transfer learning. It was later found that generative LLM such as GPT required task-specific fine-tuning. However, AI researchers showed that scaling up LLM also enabled them to solve tasks they have not been specifically fine-tuned for, so-called few-shot or even zero-shot learning [24]. These LLM achieved translation, question-answering, and cloze tasks with only seeing few or zero examples, and even accomplished reasoning-related capabilities [24,30]. With increases in model size (GPT 3 had 175B parameters), LLM became increasingly capable of performing unseen tasks, such as two- or three-digit arithmetic. Interestingly, some capabilities only emerge with sufficient model size (number of parameters), compute steps during training (e.g., training epochs), and training data size [31]. Nowadays, LLM are as flexible as to pass exams in medicine, law, or business without specific fine-tuning [32]. Also for concept inventories in physics, GPT-based LLM such as ChatGPT were found to achieve almost 100% accuracy [33,34]. Given these findings, it is the next step to find out whether and how to utilize LLM-based tools such as ChatGPT for pilot-testing purposes for CIs.

### D. Data augmentation techniques

The widespread use of artificial neural networks is mainly attributed to their exceptional performance in

certain classification and pattern recognition tasks. To achieve such high performance, the training of neural networks requires a large amount of data. However, in some cases, it is difficult to collect a large amount of data, such as empirical data in education or time-series data [35]. This is critical because a smaller number of training data usually leads to a reduced generalization ability and a lower overall performance of the model. One solution to increase the generalization ability of trained models by reducing the chance of overfitting is the generation of synthetic data, the so-called data augmentation [36].

In the field of image recognition, data augmentation is a common practice in which, for example, the original image is cropped, mirrored, or its colors are changed to enhance the number of images in the training set [37]. In time-series data, a recent literature review provides an overview of magnitude, time domain, and frequency domain transformation techniques [35]. In an additional empirical evaluation of various data augmentation methods for time-series data, the authors report both enhanced and decreased generalization accuracy. For example, a permutation method in which the observations and time of the observations are changed, leads to detrimental effects, as it breaks the time dependence of the series.

In our case, the educational data to be enhanced are the solutions on a standardized concept inventory, the FCI. The analysis of data augmentation of educational data is important for test instrument development and validation as well as for adaptive learning systems that provide a personalized response to a learner's input or behavior. In concept tests, the items are related to the concepts, but there is no inherent sequence dependence. Often the items are even administered in a random order to avoid potential effects of the item sequence. Therefore, typical data augmentation techniques in the range of time-domain transformations are not suitable for educational data. In addition, students' responses to concept test items are not three dimensional as images ( $x$  coordinate,  $y$  coordinate, color). Therefore, common techniques of image augmentation are not a promising augmentation method for educational data.

One technique that was used across different fields of data augmentation is jittering, in which a small offset is given to the original data. However, educational data such as students' responses to concept test items are discrete and not continuous, and small variations in the responses or the item, in which a certain response was given, would lead to randomly incorrect responses. Consequently, this technique would have a similar effect as permutation techniques and it is not likely to perform well.

Therefore, it is important not to translate these existing methods for educational data augmentation but to be sensitive to the inherent structure of the data and study options that consider the structure of concepts in the test items. In contrast to the data types mentioned above, educational data often exhibit a hierarchical structure, e.g., a student is in a specific class with a certain teacher

and the class is part of a specific school. To account for this, it is important that a suitable augmentation technique for educational data would consider different levels of prior knowledge. As mentioned earlier, large language models have made a tremendous advance in the past years and they have shown to be able to respond reasonably well to physics concept test items. Therefore, they are a promising method to augment educational data, but it is an open question to what extent and how systematically a large language model can reflect different levels of prior knowledge and students' difficulties.

To assess the suitability for generating a valuable synthetic dataset in the field of education and to account for varying levels of prior knowledge, this study explores the following research questions:

RQ1: What is the conceptual understanding of the force concept of ChatGPT?

RQ2: To what extent is ChatGPT able to simulate conceptual understanding of different cohorts?

RQ3: Is ChatGPT able to simulate specific preconceptions of students?

### III. METHODS

#### A. Participants

We compared the responses to the FCI given by ChatGPT with real students in the first semester at the RPTU Kaiserslautern-Landau, Germany. The engineering students attended a lecture of experimental physics covering the topic of mechanics, among others. The FCI was administered at the beginning of the semester, so at the moment of answering the test items, the students had only covered these concepts in high school but not yet during the lecture. In total, there were 121 engineering students (12 female, 105 male, and 4 made no statement) with an average age of 20.6 years.

#### B. The Force Concept Inventory

To examine to what extent ChatGPT can be used fruitfully in educational contexts for tasks such as solving standardized test instruments and generating data, we evaluated performance in the domain of mechanics as it is coherently structured with regard to knowledge. Physics education researchers devised standardized test instruments that are widely employed across universities throughout the world. Among the most well-known instruments are concept inventories that probe learners' deep understanding of a particular concept. Physics educators established that learners come with intuitive conceptions of certain physics-related concepts, also referred to as common-sense knowledge [38]. These might not coincide with more normative ideas about physics, such as Newtonian physics and are typically context-dependent [38]. Most learners developed over his or her upbringing an intuition of forces as acting agents [6,39], both by experiencing and interacting with the

world and hearing language such as “gravity pulls.” These intuitive conceptions are largely incompatible with the rich physics conceptions of forces, such as in Newtonian mechanics where Newton’s laws lay out a consistent conceptualization of forces. The first law states that an object at constant velocity is remaining in this state of constant velocity unless a force is acting upon it. The second law states that the acceleration of the object is dependent upon the force and mass of the object. Finally, the third law states that every force has an opposite equal force (*actio equals reactio*).

To probe students’ understanding of Newtonian mechanics, the FCI was developed [6], originally in 1992, and revised in 1995, containing 30 multiple-choice items. The FCI, as is typical for physics problems, combines image and text data. However, as ChatGPT based on GPT 4 only accepts text data as input, we transcribed the pictures in the FCI, similar to the procedure employed by West [34]. For example, a typical FCI item would include a picture of the following scenario:

A heavy ball is attached to a thread and moved horizontally around in a circle. At one point, the thread suddenly breaks. The whole process is viewed from above. Which trajectory best describes the path of the ball after the thread breaks?

1. a trajectory curve that nearly follows the path of the circle
2. a straight-line trajectory curve that points tangentially along the circle at the break point
3. a straight-line trajectory curve that points straight away from the circular path at an angle of  $45^\circ$  from the tangent at the break point
4. a trajectory curve that initially goes outward from the center, but then takes a trajectory curve in the direction of the original motion
5. a trajectory curve that points straight outward from the center [40].

West [34] verified that besides the transcriptions, ChatGPT is still capable of solving the FCI. The transcription of all items that contained a figure can be found in the Appendix.

### C. RQ1: Probing conceptual understanding of ChatGPT in physics

Probing the accuracy of ChatGPT on physics contents and the FCI was performed by multiple researchers [33,34,41]. Some studies examined qualitative physics-related dialogues with ChatGPT and short-form physics essays [42,43]. Gregorcic and Pendrill [42] found for a simple prompt (“A teddy bear is thrown into the air. What is its acceleration in the highest point?”) that ChatGPT’s responses (based on GPT version 3.5) were full of contradictions and concluded that ChatGPT is not good enough as a cheating tool in physics. Yeadon *et al.* [43] used ChatGPT (based on version 3.5) to generate responses (300-word essays) to questions and found that it performed on par with

the students. Plagiarism detection software was largely incapable of distinguishing it from human writing. Kortemeyer [33] then used ChatGPT (version 3.5) on the FCI and found that it scored correctly 18 out of 30 questions (60%). This would equal approximately the performance of a beginning learner. To ensure that ChatGPT was not simply trained on the FCI (or at least could generalize beyond mere textual matching), Kortemeyer changed surface-level information in an item and found that ChatGPT was not distracted by these changes and still performed correctly. It was also found that ChatGPT made several errors that related to the impetus (*I*) preconception. Moreover, other errors related to the confusion of individual forces versus net force on an object, unstable concepts, and logical errors [33]. West [34] could replicate these findings for ChatGPT based on GPT version 3.5 and found that a substantial improvement to expert-level performance (22 correct out of 23) could be observed when using ChatGPT based on GPT version 4. In RQ1, we will generate a larger sample of response sets for ChatGPT (based on GPT version 4) and calculate the accuracy and standard deviation of the responses against the correct answers. This will enable us to estimate the degree to which ChatGPT is certain of its answers. Similar to West [34], we will then compare ChatGPT’s distribution of the number of correct answers to the distribution of a student cohort.

### D. RQ2: Simulating conceptual understanding of different cohorts

An important capability of LLM for education is to simulate (synthesize) data, given specific prompting strategies. In general, pretrained LLM generate output based on the context provided in the input. The provided input context determines the output, given that a different sequence of tokens is to be predicted. It is then possible to adjust the response behavior of ChatGPT for the FCI questions. For example, one could request that ChatGPT always answers randomly for each question. A more informative prompt would be to ask ChatGPT to answer from the perspective of a different student group, which could be related to background variables, such as experience. In RQ2, we chose the experience-related variable and prompted ChatGPT to generate responses from the perspective of (a) school students ( $N = 8$ ) and (b) engineering students ( $N = 18$ ). We chose these groups to potentially generate variance and see to what extent differences emerge in ChatGPT’s responses. We then generate for each cohort a number of response sets in order to analyze the distribution of correct answers, comparing the cohorts with the real students and the plain ChatGPT response sets. As a robust (nonparametric) test statistic, suited for multiple groups, we utilized the Kruskal-Wallis rank test [44]. This test does not assume the normality of residuals, though it assumes identically shaped and scaled distributions for all groups,

and ordinal dependent variables. Effect size (proportion of variance predicted from the independent variable) was calculated as  $(H - k + 1)/(N - k)$ , where  $H$  is the Kruskal-Wallis test statistic,  $k$  is the number of groups (here 3), and  $N$  is the total number of observations. This effect size is called  $\eta^2$ , where  $\eta^2 = 0.01$  is a small effect size,  $\eta^2 = 0.06$  is a medium effect size, and  $\eta^2 = 0.14$  is a large effect size [45].

### E. RQ3: Simulating students' preconceptions with ChatGPT

Similar to simulating different student cohorts, it is equally possible to simulate students with different preconceptions. The FCI is a considerably designed instrument, based on an assessment of students' beliefs about force against a normative theory, namely Newtonian physics. It is recognized that overcoming preconceptions is among the crucial tasks in physics and science education, and it is a hard problem. The FCI can help instructors to align their instruction and assess students' understanding. The FCI addresses different preconceptions that students can hold. The authors identify general categories of preconception, namely kinematics (K), impetus (I), active force (AF), action-reaction pairs (AR), concatenation of influences (CI), and other influences on motion (OI). Category K indicates if learners have a physics notion of motion description. Category I refers to learner's notion that force is some kind of intrinsic quantity to objects that keeps them moving, once it has been supplied to them. AF then refers to the conception that motion implies an active force. AR refers to the learner's conception of interactions as a struggle where the stronger or more forceful party exerts a greater force (most typically exemplified in the crash of a truck and a car). CI aggregates several preconceptions, such as the dominance principle in the composition of two unequal forces on a single object. Finally, also OI subsumes several preconceptions such as the reasoning with centrifugal force and considering it as a real force.

Given the well-established theory on preconceptions related to the Newtonian mechanics and forces in particular, we can utilize this information to prompt ChatGPT to generate (simulate) responses as if it were a learner with this preconceptions. To do so, we deliberately engineered the prompts given to ChatGPT. We first checked if ChatGPT knew about (could elaborate on) the respective preconception (K, I, AF, AR, CI, OI). In fact, ChatGPT knew about all preconceptions and could elaborate on what they were. Afterward, we engineered a prompt where the items were preceded with the following phrase: "Respond to the following questions from the perspective of a person who has the preconception that ...". We then described the preconception in brief or simply inserted "the impetus preconception" for agreed-upon preconception [33]. For each preconception, at least seven different response sets

(more responses were generated for preconceptions with higher observed variability in the responses) were generated. We then analyze the distribution of correct answers. Furthermore, we prompted ChatGPT for important items to provide reasoning for their responses, given that it simulated a student with a different preconception. We validated if the reasoning was in alignment with the preconception.

## IV. FINDINGS

### A. RQ1: Probing conceptual understanding of ChatGPT in physics

Regarding RQ1, we first transformed the response matrix into a binary coded matrix where correct responses were coded with 1 and incorrect responses were coded with 0. Based on this binary matrix, we calculated the average correct answers (of overall  $N = 14$  responses) and the standard deviation. We found that  $M = 0.83$  ( $SD = 0.22$ ), i.e., ChatGPT scored on average 83% in the FCI (with some transcribed items). In Table I, the distributional statistics for all items with a mean score below 1.0 is displayed. It can be verified that for many items, ChatGPT perfectly solved them. With other items (e.g., A8), there is quite some variance, and solution probability is almost equal to random guessing. Item A20 was never solved correctly by ChatGPT. The engineering students ( $N = 121$ ) performed on average  $M = 0.14$  ( $SD = 0.20$ ).

### B. RQ2: Simulating conceptual understanding of different cohorts

Next, we prompted ChatGPT to answer the FCI as if it were (a) a 10th-grade student in school or (b) an engineering student. In Table II, the means and standard deviations for the distributions of the groups are depicted. In a reasonable approximation, we can verify that the scales (i.e., standard deviations) are similar for the groups, so we meet the assumptions for the Kruskal-Wallis test. This can also be visually inspected in Fig. 1, where the distributions of correct responses for the different cohorts are depicted.

TABLE I. Item statistics (generated with the assistance of ChatGPT). Only items with a mean score below 1.00 are included.

| Item | Mean | Median | SD   | Variance | 5%   | 95%  |
|------|------|--------|------|----------|------|------|
| A8   | 0.14 | 0.00   | 0.36 | 0.13     | 0.00 | 1.00 |
| A13  | 0.93 | 1.00   | 0.27 | 0.07     | 0.65 | 1.00 |
| A14  | 0.79 | 1.00   | 0.43 | 0.18     | 0.00 | 1.00 |
| A16  | 0.50 | 0.50   | 0.52 | 0.27     | 0.00 | 1.00 |
| A17  | 0.71 | 1.00   | 0.47 | 0.22     | 0.00 | 1.00 |
| A18  | 0.64 | 1.00   | 0.50 | 0.25     | 0.00 | 1.00 |
| A19  | 0.36 | 0.00   | 0.50 | 0.25     | 0.00 | 1.00 |
| A20  | 0.00 | 0.00   | 0.00 | 0.00     | 0.00 | 0.00 |
| A21  | 0.57 | 1.00   | 0.51 | 0.26     | 0.00 | 1.00 |
| A29  | 0.57 | 1.00   | 0.51 | 0.26     | 0.00 | 1.00 |
| A30  | 0.71 | 1.00   | 0.47 | 0.22     | 0.00 | 1.00 |

TABLE II. Descriptive statistics for groups.

| Group       | Mean  | Median | SD   | Variance | 5%   | 95%  |
|-------------|-------|--------|------|----------|------|------|
| Engineers   | 24.28 | 25.0   | 1.02 | 1.04     | 22.0 | 25.0 |
| Unspecified | 24.93 | 25.0   | 1.27 | 1.61     | 23.0 | 27.0 |
| Students    | 24.25 | 24.5   | 1.58 | 2.50     | 22.0 | 26.0 |

The Kruskal-Wallis test yields a value for potential group differences of  $H = 1.79$ ,  $p = 0.409$ ,  $\eta^2 = -0.01$ . Essentially, no differences were found between the aggregated correct response distributions for the three cohorts.

### C. RQ3: Simulating students' preconceptions with ChatGPT

Figures 2 and 3 show the distributions for the simulated students with different preconceptions. It can be seen that they are widely spread and depending on the specific preconceptions more or less accuracy can be observed in the response sets. Interestingly, students with the AR preconception almost systematically fail at all items, whereas simulated students with CF solve the FCI as accurately as without preconceptions 3). Certainly, a much greater variability as in randomly generated responses can be observed.

Finally, we qualitatively validated two preconceptions with regard to the provided reasoning by ChatGPT. As for their intricate and prevalent nature, we chose centrifugal force (CF) and impetus (I). The centrifugal force is not a real force as for example gravity, but rather a pseudoforce that results from the chosen frame of reference. In the inertial frame of reference, only the centripetal force acts on the object attached to the string that keeps it on the circular track. Once the string is cut, no force (except for gravity and normal force) is acting on the moving object, which, according to Newton's first law keeps on moving with

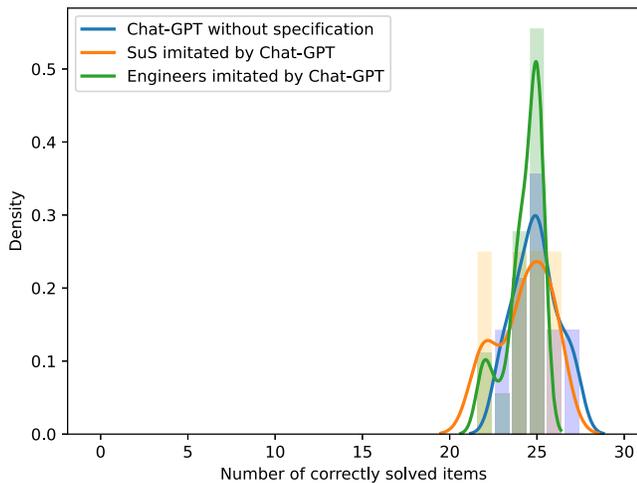


FIG. 1. Distribution of correct responses for different cohorts as prompted and generated by ChatGPT.

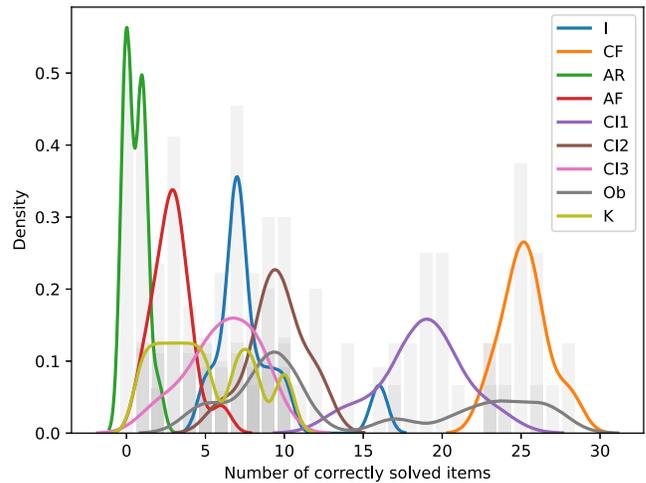


FIG. 2. Distribution of correct responses for different preconceptions as prompted and generated by ChatGPT.

constant velocity (assuming frictionless movement). The centrifugal force during a circular motion in a stationary frame of reference (also termed reactive centrifugal force, e.g., in a car) feels very real for students as they experience this force [46]. Centrifugal force is assumed in a center of mass frame of reference that moves alongside the object in the circular track. No net force should act on the object, as in this frame of reference, it is at rest. Students with this preconception would choose answer 4, i.e., after the string is plugged, the outward force (centrifugal force) will drive the object outwards. ChatGPT reasons this case convincingly and correctly. It is noteworthy that no mention of the frame of reference problem is made.

**Prompt**, (see Sec. III):

**Responses by ChatGPT (translated by ChatGPT):** First response: From the perspective of a person who incorrectly believes that a

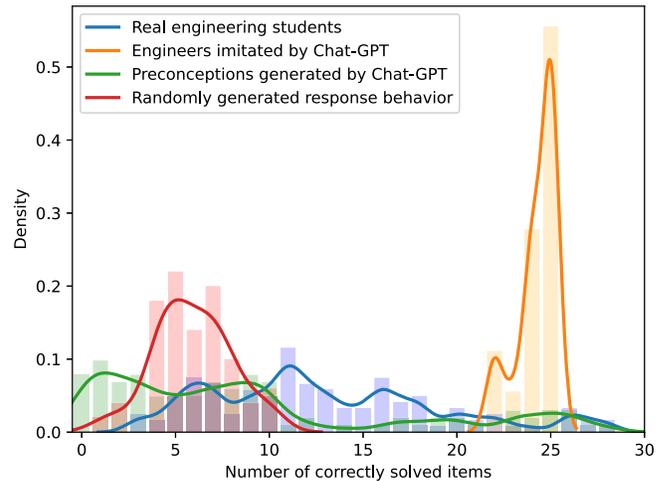


FIG. 3. Distribution of correct responses for different cohorts as prompted and generated by ChatGPT.

centrifugal force exists, they may choose Answer 4: “A trajectory that initially goes from the center outwards, but then takes a trajectory in the direction of the original movement.”

The misconception of centrifugal force is based on the feeling that a force is acting during circular motion that pushes the body outwards. In reality, it is the centripetal force (a force directed inwards) that keeps the body on its circular path. This misconception could lead a person to believe that when the string breaks, the ball is initially driven outwards from the center (by the supposed centrifugal force) before then moving in the direction of the original motion because there is no longer any centrifugal force to push it outwards.

The correct answer would, however, be Answer 2: “A linear trajectory that shows tangential along the circle at the breaking point.” Because when the string breaks, there is no longer a centripetal force that forces the ball into a circular path, and the ball continues its motion in a straight line, which is tangential to the original circular path, due to inertia. This motion corresponds to Newton’s first law, which states that a body remains at rest or moves with constant velocity in a straight line unless a force acts on it.

As for the impetus concept, students typically believe (much in agreement with their lived experiences in a friction-based world) that when an object is thrown, something is transferred to the object that it can use to fly a certain time, constantly using up the impetus. Once it is used up, the objects come to rest. We would expect students to choose answer 1 for this item. After all, a circular impetus is given to the object (e.g., by the string). Once the string is plugged, this circular impetus will cause the object to continue on the circular path. Again, ChatGPT correctly reasoned about this preconception and chose the correct answer, responding as if it were a student with the impetus preconception.

**Prompt**, (see Sec. III):

**Responses by ChatGPT (translated by ChatGPT):** Second response: The theory of impetus is a medieval physical theory that was proposed before Newton’s laws of motion. This theory suggests that a moving object possesses an “impetus” that drives it further in the direction of the original motion, even when no external force acts on it anymore. It is believed that the impetus gradually decreases, and the object eventually stops.

With this misconception, a person might choose Answer 1: “A trajectory that closely follows the course of the circular path.” This person might

believe that after the string breaks, the ball continues to fly in the original circular motion because the impetus drives it in that direction. The ball would then gradually deviate from the circular path and finally stop when the impetus has completely diminished.

However, the correct answer according to modern laws of physics would be answer 2: “A straight-line trajectory that is tangential to the circle at the breaking point.” According to Newton’s first law of motion, in the absence of an external force (which was removed by the breaking string), the ball would continue its motion in a straight line, which is tangential to the circular path it was originally on.

## V. DISCUSSION

The recent advances in AI research, in particular generative AI, provide novel potentials for physics education researchers to enhance research and teaching, however, they also introduce novel challenges. To better understand the potentials and challenges, it is crucial to engage with advanced technologies, such as LLM, to evaluate their performance for physics education-specific problems (as argued by Wang in Ref. [47], who called for “more careful and controlled studies” on this matter, particularly with ChatGPT). In this study, we evaluated the capabilities of an LLM, namely ChatGPT, to generate synthetic data for the FCI and thus potentially enhance concept inventory development. We chose the FCI as a use case as this is a well validated and most widely used concept inventory in physics. Prior studies established that ChatGPT is well versed in solving the FCI. In our study, we advanced this research to compare performance with real students (RQ1) and generate synthetic data for it with regard to different cohorts (RQ2) and varying preconceptions (RQ3).

In RQ1, we found that ChatGPT (based on GPT 4) solved the FCI (with transcribed images) accurately. In fact, it solved the FCI much more accurately compared to real engineering students in a German university. This resonates with prior studies that established that ChatGPT can accurately solve the FCI [33,34]. Compared to the prior work by West, where ChatGPT (based on GPT-4) solved 22 out of 23 items correctly, here ChatGPT (based on GPT-4) only correctly solves about 26 out of 30 items on average. There were 11 items that had not been solved correctly in each run (see Table I). Of these 11 items, 6 contained a figure that was translated into text. In comparison, four items that contained text translated from a figure were solved correctly in each run. Consequently, the translation from figures to text may cause difficulties for ChatGPT to interpret. Previously, some of these items were also taken out from consideration [34]. As a note of interest, ChatGPT can be easily prompted to output the responses in varying formats (e.g., comma-separated list), and it solves it much

more quickly than real students. We particularly generated more response sets in order to retrieve information on the variance displayed in responding to the FCI. For most items, no variance was found, because they were solved with 100% accuracy. For other items, however, a mere random performance was observed.

In RQ2, we sought to utilize ChatGPT (via specific prompting) to simulate answers from students in different cohorts (10th-grade students and engineering students). To our surprise, we found no differences between these cohorts and the unprompted (with regard to cohort) responses. Either the mean accuracy was similar or the standard deviation was similar. This struck us as unexpected, given that it is established that prompting can in fact be used to modify the responses of ChatGPT and other generative AI-based LLM [48]. However, maybe the provided prompts were not specific enough to affect the responses in ChatGPT.

Therefore, in RQ3, we specifically prompted ChatGPT to respond to the FCI as if it were a student with a certain preconception related to Newtonian force concept and mechanics more generally. Here, we found large differences in the response behavior of ChatGPT. ChatGPT knew about all preconceptions. The accuracy, given that ChatGPT responds as if it were a student with almost any preconception, dropped considerably, sometimes even below random guessing performance. Interestingly, prompting ChatGPT with preconceptions resulted in distributions that were equally spread as real students' distributions. For two preconceptions, we verified that ChatGPT also correctly reasoned if it was prompted to respond to the items as if it had a certain preconception, which it did.

These data augmentation techniques could be of interest to physics education researchers and educational researchers more generally if they want to pilot test their instrument, e.g., in the domain of mechanics. Even practitioners designing classroom exams and policy agents designing statewide exams could benefit from these capabilities of generative AI (here: LLM) to ensure the validity of their instruments. One caveat would be that users currently will have to outsource their instruments to a private company. On the other hand, it can be tiresome for students to function as "lab animals" and be subjected to pilot testing of research instruments (let alone ethical concerns). Moreover, the precious time of the students is consumed which could be used in more learning-relevant activities.

## VI. LIMITATIONS

While LLM such as GPT 4 provide novel capabilities for research such as data augmentation, they also introduce novel challenges that researchers have to engage with. Brown *et al.* [24] list three: misuse, bias, and energy consumption as important points of concern. For once, LLM such as GPT 4 can be used to generate human-sounding text and therefore be used to spread certain narratives. With regard to physics education research, students might fabricate their essays with

ChatGPT, which cannot be easily detected by human instructors or researchers [43]. In the future, the selection of teaching materials should be carefully tailored to ensure that tools like ChatGPT are not undesirably misused as aids. For example, we can observe a similar context where students increasingly resort to aids such as YouTube videos when the provided materials exhibit deficiencies [49]. Equally concerning is the problem of hallucination in LLM such as GPT 4. For example, ChatGPT provides well-sounding output with sometimes flawed physics [42]. This becomes a problem when considering that students accept the outputs of Chat-GPT without critical reflection (copy and paste it) [50]. Instructors need to find means to prevent tasks where simple copy paste is sufficient. We also found that important steps in the problem-solving process such as explicating underlying assumptions and idealizations are oftentimes omitted [51]. This makes it difficult for students to truly learn the underlying physics. Kortemeyer [33] prompted ChatGPT on its ability to understand physics and received the following response:

“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.”

As such, human expert judgments will remain crucially important to verify responses given by ChatGPT. As a matter of fact, we found the theory-based, specific prompting with regard to established preconceptions yielded the most variance which could be important for data augmentation. As such, established theory is crucial to guide prompt engineering and generate meaningful synthetic data. It was furthermore established that LLM such as GPT 3 and others incorporate humanlike biases, e.g., with regard to gender [24,52]. Given the fact that they are trained on the Common Crawl of the Internet and similar unclean data, this is probably not surprising. However, approaches to mitigate this problem are still not fully developed, and thus employing ChatGPT and other LLM in educational settings should probably be postponed, especially in contexts of high-stakes testing. Finally, the energy consumption of training these LLM and even requesting a single output are ecologically challenging [53]. It will likely be not sustainable for every research group to run their own models. Shared infrastructures will be required to make these applications sustainable.

## VII. IMPLICATIONS

Our study showed that with adequate prompting strategies, ChatGPT can be utilized to generate synthetic data

related to a concept inventory with meaningful variability in the dataset. We submit that similar conclusions will likely be true for other concept inventories as well, such as climate change, electricity, and biological evolution. Further research is needed to test these predictions. This offers researchers a valuable tool to simulate field testing their concept inventories and evaluate psychometric properties, given sufficient sample sizes. We could show that similar distributions as with real students emerge, and further research could investigate if convolutions of preconceptions will enable a reconstruction of the true students' distribution. We cannot, however, exclude the possibility that real students also show other response patterns, i.e., have different conceptions that are not included in our theoretically derived set of preconceptions. Further research should investigate also the possibility to even generate a concept inventory in the first place with the help of LLM such as GPT 4 [54] and assist in other phases and steps of the development process for a concept inventory [15].

The code for analysis (and further analyses) can be accessed here at Ref. [55].

The synthetic data can be accessed upon reasonable request to the authors.

This research was supported by LMUexcellent, funded by the Federal Ministry of Education and Research (BMBF) and the Free State of Bavaria under the Excellence Strategy of the Federal Government and the Länder.

### APPENDIX: TRANSLATION OF FIGURES IN FCI ITEMS

In this section, we show the translation of figures in FCI items into text so that ChatGPT is able to interpret it. We did not include the original figures due to copyright restrictions, but the item numbering is identical to the original work [6].

#### Item 5

Which of the following forces act(s) on the ball when it is exactly halfway through a semicircular channel?

- A) A downward force of gravity.
- B) A force exerted by the channel that is directed from the point where the ball is currently located to the center of the channel.
- C) A force in the direction of motion.
- D) A force directed from the center point to the point in the channel where the ball is currently located.

1. force A only.
2. forces A and B.
3. forces A and C.
4. forces A, B, and C.
5. forces A, C, and D.

#### Item 6

Which path would a ball most closely follow after it exits the channel at the end and moves across the frictionless table top?

1. A path that would correspond to the arc of the gutter if it formed a closed circle.
2. A straight-line trajectory curve that points tangentially along the circle at the exit point.
3. A trajectory curve that describes an arc of a circle outwards.
4. A path curve which describes a parabola to the outside.
5. A trajectory in which the ball is immediately deflected outward in a straight line.

#### Item 7

A steel ball is attached to a string and is swung in a circular path in a horizontal plane. At one point, the string suddenly breaks near the ball. If these events are observed from directly above, which path would the ball most closely follow after the string breaks?

1. A trajectory that almost corresponds to the course of the circular path.
2. A straight-line trajectory that points tangentially along the circle at the tear-off point.
3. A straight-line trajectory that points away from the circular path at an angle of  $45^\circ$  from the tangent line at the break-off point.
4. A trajectory that initially goes outward from the center but then takes a trajectory in the direction of the original motion.
5. A trajectory that points outward in a straight line from the center.

#### Item 8

Imagine a hockey puck is sliding with constant speed  $v_0$  in a straight line from point "P" to point "Q" on a frictionless horizontal surface. Forces exerted by the air are negligible. You are looking down on the puck. When the puck reaches point "Q," it receives a swift horizontal kick perpendicular to the previous direction of motion. Had the puck been at rest at point "Q," then the kick would have set the puck in horizontal motion with a speed  $v_k$  in the direction of the kick.

Which paths would the puck most closely follow after receiving the kick?

1. A straight line in the direction of the collision.
2. A straight line pointing at a sharp angle to the previous trajectory.
3. A path in the same direction as the original direction of motion but offset in parallel in the direction of the impact.
4. A curved path that slopes in the direction of the shock.
5. A curved path inclined in the direction of the original direction of motion.

**Item 12**

A ball is fired by a cannon from the top of a cliff in the horizontal direction. Which paths would the cannon ball most closely follow?

1. A path that runs in a straight line down to the point of impact.
2. A path that is curved toward the point of impact.
3. A path that is curved toward the point of impact, but which is horizontal for a short distance after launch.
4. A path that is curved toward the point of impact, but which is horizontal for a certain distance after launch and vertical before impact.
5. A path that runs horizontally until almost vertically above the point of impact and then runs vertically until the point of impact.

**Item 14**

A bowling ball accidentally falls out of the cargo bay of an airliner as it flies along in a horizontal direction. As observed by a person standing on the ground and viewing the plane from the side, which path would the bowling ball most closely follow after leaving the airplane?

1. A curved path that runs backward.
2. A path that runs vertically downward.
3. A path that runs forward in a straight line.
4. A curved path that runs forward.
5. A path that extends forward horizontally to above the point of impact and then vertically downward.

**Items 15, 16 and 17**

The figures have not been described as they only carry redundant information that is also contained in the text.

**Item 18** A child swings on a rope tied to a branch (point O). His movement starts above a point P and then continues downward through the point P. Consider the following distinct forces:

1. A downward force of gravity.
2. A force exerted by the rope pointing from A to O.
3. A force in the direction of the boy's motion.
4. A force pointing from O to A.

Which of the above forces is (are) acting on the boy when he is at position A?

- (A) 1 only.
- (B) 1 and 2.
- (C) 1 and 3.
- (D) 1, 2, and 3.
- (E) 1, 3, and 4.

**Item 19**

Two blocks move to the right. The following table shows the distance covered by the blocks at intervals of 0.2 s in length units (LU). Identical numbers indicate identical points in time.

Block a: 1: 2 LU, 2: 4 LU, 3: 7 LU, 4: 11 LU, 5: 16 LU, 6: 22 LU, 7: 29 LU

Block b: 1: 0 LU, 2: 4 LU, 3: 8 LU, 4: 12 LU, 5: 16 LU, 6: 20 LU, 7: 24 LU, 8: 28 LU

Do the blocks ever have the same speed?

1. No.
2. Yes, at instant 2.
3. Yes, at instant 5.
4. Yes, at instants 2 and 5.
5. Yes, at some time during the interval 3 to 4.

**Item 20**

Two blocks move to the right. The following table shows the distance covered by the blocks at intervals of 0.2 s in length units (LU). Identical numbers indicate identical points in time.

Block a: 1: 2 LU, 2: 6 LU, 3: 10 LU, 4: 14 LU, 5: 18 LU, 6: 22 LU, 7: 26 LU

Block b: 1: 0 LU, 2: 6 LU, 3: 12 LU, 4: 18 LU, 5: 24 LU

The accelerations of the blocks are related as follows:

- (A) The acceleration of "a" is greater than the acceleration of "b".
- (B) The acceleration of "a" equals the acceleration of "b." Both accelerations are greater than zero.
- (C) The acceleration of "b" is greater than the acceleration of "a."
- (D) The acceleration of "a" equals the acceleration of "b." Both accelerations are zero.
- (E) Not enough information is given to answer the question.

**Item 21**

A rocket drifts sideways in outer space from point "a" to point "b" as shown below. The rocket is subject to no outside forces. Starting at position "b," the rocket's engine is turned on and produces a constant thrust (force on the rocket) at right angles to the line "ab." The constant thrust is maintained until the rocket reaches a point "c" in space.

Which path below best represents the path of the rocket between points "b" and "c"?

1. Point "c" lies to the right above point "b." The path from "b" to "c" is concave.
2. Point "c" is right above point "b." The path from "b" to "c" is straight and perpendicular to the line connecting "b" to "a."
3. Point "c" lies to the right above point "b." The path from "b" to "c" is straight.
4. Point "c" lies above and to the right of point "b." The path curve from "b" to "c" runs horizontally to the right at the beginning and is convex from then on.
5. Point "c" lies above and to the right of point "b." The path curve from "b" to "c" is convex.

**Item 23**

At point "c," the rocket's engine is turned off and the thrust immediately drops to zero.

Which of the paths below will the rocket follow beyond point "c"?

1. The path curve runs in a straight line horizontally to the right.
2. The path curve runs in a straight line diagonally upwards to the right.

3. The path curve runs vertically upwards in a straight line.
4. The path curve runs to the right and is concave.
5. The path curve runs to the right and is convex.

**Item 28** The figure has not been described as it only carries redundant information that is also contained in the text.

- 
- [1] E. Kasneci, K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günemann, E. Hüllermeier *et al.*, ChatGPT for good? On opportunities and challenges of large language models for education, *Learn. Individ. Diff.* **103**, 102274 (2023).
- [2] R. Van De Schoot, J. De Bruin, R. Schram, P. Zahedi, J. De Boer, F. Weijdem, B. Kramer, M. Huijts, M. Hoogerwerf, G. Ferdinands *et al.*, An open source machine learning framework for efficient and transparent systematic reviews, *Nat. Mach. Intell.* **3**, 125 (2021).
- [3] T. Kuzman, I. Mozetič, and N. Ljubešić, ChatGPT: Beginning of an end of manual linguistic data annotation? Use case of automatic genre identification, [arXiv:2303.03953](https://arxiv.org/abs/2303.03953).
- [4] N. M. S. Surameery and M. Y. Shakor, Use ChatGPT to solve programming bugs, *Int. J. Inf. Technol. Comput. Eng.* **3**, 17 (2023).
- [5] R. K. Sinha, A. Deb Roy, N. Kumar, and H. Mondal, Applicability of ChatGPT in assisting to solve higher order problems in pathology, *Cureus* **15**, e35237 (2023).
- [6] D. Hestenes, M. Wells, and G. Swackhamer, Force concept inventory, *Phys. Teach.* **30**, 141 (1992).
- [7] I. T. Koponen and M. Pehkonen, Coherent knowledge structures of physics represented as concept networks in teacher education, *Sci. Educ.* **19**, 259 (2010).
- [8] P. Wulff, Network analysis of terms in the natural sciences insights from Wikipedia through natural language processing and network analysis, *Educ. Inf. Technol.* (2023), [10.1007/s10639-022-11531-5](https://doi.org/10.1007/s10639-022-11531-5).
- [9] X. Liu, *Using and Developing Measurement Instruments in Science Education: A Rasch Modeling Approach*, Science & Engineering Education Sources (Information Age Publishing Incorporated, Charlotte, NC, 2010).
- [10] L. Porter, C. Taylor, and K. C. Webb, Leveraging open source principles for flexible concept inventory development, in *Proceedings of the 2014 Conference on Innovation & Technology in Computer Science Education*, edited by Å. Cajander, M. Daniels, T. Clear, and A. Pears (ACM, New York, NY, 2014), pp. 243–248.
- [11] R. H. Nehm and H. Härtig, Human vs. computer diagnosis of students' natural selection knowledge: Testing the efficacy of text analytic software, *J. Sci. Educ. Technol.* **21**, 56 (2012).
- [12] T. Schubatzky, R. Wackermann, C. Wöhlke, C. Haagen-Schützenhöfer, M. Jedamski, H. Lindemann, and K. Cardinal, Entwicklung des concept-inventory ccci-422 zu den naturwissenschaftlichen grundlagen des klimawandels, *Z. Didakt. Nat. Wiss.* **29**, 10 (2023).
- [13] 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).
- [14] P. Eaton, Evidence of measurement invariance across gender for the force concept inventory, *Phys. Rev. Phys. Educ. Res.* **17**, 010130 (2021).
- [15] W. K. Adams and C. E. Wieman, Development and validation of instruments to measure learning of expert-like thinking, *Int. J. Sci. Educ.* **33**, 1289 (2011).
- [16] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, A neural probabilistic language model, *J. Mach. Learn. Res.* **3**, 1137 (2003).
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez *et al.*, Attention is all you need: Conference on neural information processing systems, in *Advances in Neural Information Processing Systems (Bd. 17)*, edited by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan *et al.* (Curran Associates, Inc., 2017), <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- [18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [19] C. D. Manning, Human language understanding & reasoning, *Daedalus (Boston)* **151**, 127 (2022).
- [20] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, Scaling laws for neural language models (2020).
- [21] S. Wolfram, *What is ChatGPT Doing and Why Does it Work?* (Wolfram Media, 2023).
- [22] More precisely, tokens are used instead of words, which are more general sequences of letters that can be added to form words. However, they have the advantage that also unseen words in the training data can be processed.
- [23] Percentages are taken from <https://en.wikipedia.org/wiki/GPT-3#GPT-3.5> (last access 26 May 2023) for training GPT 3.
- [24] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, Language models are few-shot learners, *Adv. Neural Inf. Process. Syst.* **33**, 1877 (2020).
- [25] R. Rosenbaum, On the relationship between predictive coding and backpropagation, *PLoS One* **17**, e0266102 (2022).
- [26] R. Bommasani *et al.*, On the opportunities and risks of foundation models (2022).
- [27] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in *Proceedings of the*

- 2021 ACM Conference on Fairness, Accountability, and Transparency (2021), pp. 610–623.
- [28] A. Lewkowycz, A. Andreassen, D. Dohan, E. Dyer, H. Michalewski, V. Ramasesh, A. Slone, C. Anil, I. Schlag, T. Gutman-Solo, Y. Wu, B. Neyshabur, G. Gur-Ari, and V. Misra, Solving quantitative reasoning problems with language models (2022).
- [29] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, Large language models are zero-shot reasoners, *Adv. Neural Inf. Process. Syst.* **35**, 22199 (2022), [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/8bb0d291acd4acf06ef112099c16f326-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/8bb0d291acd4acf06ef112099c16f326-Paper-Conference.pdf).
- [30] J. Huang and K. C.-C. Chang, Towards reasoning in large language models: A survey (2023).
- [31] S. R. Bowman, Eight things to know about large language models (2023).
- [32] D. M. Katz, M. J. Bommarito, S. Gao, and P. D. Arredondo, GPT-4 passes the bar exam, [10.2139/ssrn.4389233](https://ssrn.com/abstract=4389233).
- [33] G. Kortemeyer, Could an artificial-intelligence agent pass an introductory physics course?, *Phys. Rev. Phys. Educ. Res.* **19**, 010132 (2023).
- [34] C. G. West, AI and the FCI: Can ChatGPT project an understanding of introductory physics?, [arXiv:2303.01067](https://arxiv.org/abs/2303.01067).
- [35] B. K. Iwana and S. Uchida, An empirical survey of data augmentation for time series classification with neural networks, *PLoS One* **16**, e0254841 (2021).
- [36] D. A. Van Dyk and X.-L. Meng, The art of data augmentation, *J. Comput. Graph. Stat.* **10**, 1 (2001).
- [37] C. Shorten and T. M. Khoshgoftaar, A survey on image data augmentation for deep learning, *J. Big Data* **6**, 60 (2019).
- [38] D. Hammer, Student resources for learning introductory physics, *Am. J. Phys.* **68**, S52 (2000).
- [39] A. diSessa, Toward an epistemology of physics, *Cognit. Instr.* **10**, 105 (1993).
- [40] Chat can be accessed here: <https://chat.openai.com/share/8ae109d2-8b71-4033-8c62-749c33c2582c>, last access: June 2023; translated with: <https://www.deepl.com/de/translator/#de/en/>.
- [41] C. G. West, Advances in apparent conceptual physics reasoning in ChatGPT-4, [arXiv:2303.17012](https://arxiv.org/abs/2303.17012).
- [42] B. Gregorcic and A.-M. Pendrill, ChatGPT and the frustrated socrates, *Phys. Educ.* **58**, 035021 (2023).
- [43] W. Yeadon, O.-O. Inyang, A. Mizouri, A. Peach, and C. P. Testrow, The death of the short-form physics essay in the coming AI revolution, *Phys. Educ.* **58**, 035027 (2023).
- [44] W. H. Kruskal and W. A. Wallis, Use of ranks in one-criterion variance analysis, *J. Am. Stat. Assoc.* **47**, 583 (1952).
- [45] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences* (Erlbaum, Hillsdale, NJ, 1988).
- [46] S. Küchemann, P. Klein, H. Fouckhardt, S. Gröber, and J. Kuhn, Students' understanding of non-inertial frames of reference, *Phys. Rev. Phys. Educ. Res.* **16**, 010112 (2020).
- [47] J. Wang, ChatGPT: A test drive, *Am. J. Phys.* **91**, 255 (2023).
- [48] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing (2021).
- [49] C. Ruggieri, Students' use and perception of textbooks and online resources in introductory physics, *Phys. Rev. Phys. Educ. Res.* **16**, 020123 (2020).
- [50] L. Krupp, S. Steinert, M. Kiefer-Emmanouilidis, K. E. Avila, P. Lukowicz, J. Kuhn, S. Küchemann, and J. Karolus, Unreflected acceptance – investigating the negative consequences of chatGPT-assisted problem solving in physics education, [arXiv:2309.03087](https://arxiv.org/abs/2309.03087).
- [51] F. Kieser and P. Wulff, Using large language models to probe cognitive constructs, augment data, and design instructional materials, in *Machine Learning in Educational Sciences: Approaches, Applications and Advances*, edited by M. S. Khine (Springer Nature, to be published).
- [52] A. Caliskan, J. J. Bryson, and A. Narayanan, Semantics derived automatically from language corpora contain human-like biases, *Science* **356**, 183 (2017).
- [53] J. Dodge, T. Prewitt, R. Tachet Des Combes, E. Odmark, R. Schwartz, E. Strubell, A. S. Luccioni, N. A. Smith, N. DeCario, and W. Buchanan, Measuring the carbon intensity of AI in cloud instances, in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (2022), pp. 1877–1894.
- [54] S. Küchemann, S. Steinert, N. Revenga, M. Schweinberger, Y. Dinc, K. E. Avila, and J. Kuhn, Can chatGPT support prospective teachers in physics task development?, *Phys. Rev. Phys. Educ. Res.* **19**, 020128 (2023).
- [55] [https://colab.research.google.com/drive/1WU\\_cBfcYMO2-dxE9kFVjK9Hlq0mSati8V?usp=sharing](https://colab.research.google.com/drive/1WU_cBfcYMO2-dxE9kFVjK9Hlq0mSati8V?usp=sharing).