
Probing the Critical Point (CritPt) of AI Reasoning: a Frontier Physics Research Benchmark

Minhui Zhu^{1*}, Minyang Tian^{1,2*}, Xiaocheng Yang², Tianci Zhou³, Lifan Yuan²,
Penghao Zhu⁴, Eli Chertkov⁵, Shengyan Liu², Yufeng Du², Ziming Ji⁶, Indranil Das²,
Junyi Cao², Yufeng Du⁷, Jiabin Yu⁸, Peixue Wu⁹, Jinchen He^{1,10}, Yifan Su¹¹, Yikun Jiang⁶,
Yujie Zhang^{12,9}, Chang Liu¹³, Ze-Min Huang¹⁴, Weizhen Jia¹⁵, Yunkai Wang^{12,9},
Farshid Jafarpour¹⁶, Yong Zhao¹, Xinan Chen², Jessie Shelton², Aaron W. Young¹⁷,
John Bartolotta⁵, Wenchao Xu^{18,19}, Yue Sun²⁰, Anjun Chu²¹, Victor Colussi⁵,
Chris Akers²², Nathan Brooks²³, Wenbo Fu², Jinchao Zhao²⁴, Marvin Qi²¹, Anqi Mu¹¹,
Yubo Yang²⁵, Allen Zang²¹, Yang Lyu²⁶, Peizhi Mai², Christopher Wilson²², Xuefei Guo²,
Juntai Zhou², Daniel Inafuku², Chi Xue⁵, Luyu Gao²⁷, Ze Yang², Yair Hein¹⁶,
Yonatan Kahn^{28,29}, Kevin Zhou²⁶, Di Luo³⁰, John Drew Wilson²², Jarrod T. Reilly²²,
Dmytro Bandak⁵, Ofir Press³¹, Liang Yang³², Xueying Wang³³, Hao Tong²,
Nicolas Chia¹, Eliu Huerta^{1,2,21}, Hao Peng²

¹Argonne National Laboratory ²University of Illinois Urbana-Champaign ³Virginia Tech
⁴Ohio State University ⁵Independent ⁶Northeastern University ⁷Caltech
⁸University of Florida ⁹University of Waterloo ¹⁰University of Maryland, College Park
¹¹Columbia University ¹²Perimeter Institute for Theoretical Physics
¹³University of Connecticut ¹⁴University of Cologne ¹⁵The Chinese University of Hong Kong
¹⁶Utrecht University ¹⁷Harvard University ¹⁸ETH Zürich ¹⁹Paul Scherrer Institute
²⁰University of Washington Seattle ²¹University of Chicago ²²University of Colorado Boulder
²³Chi 3 Optics ²⁴Hong Kong University of Science and Technology ²⁵Hofstra University
²⁶University of California, Berkeley ²⁷Carnegie Mellon University
²⁸University of Toronto ²⁹Vector Institute ³⁰University of California, Los Angeles
³¹Princeton University ³²University of California San Diego
³³National Institute of Theory and Mathematics in Biology

* Equal contribution lead authors.

Correspondence to minhui.zhu@anl.gov, mtian8@illinois.edu.

 critpt.com  CritPt  CritPt

Abstract

While large language models (LLMs) with reasoning capabilities are progressing rapidly on high-school math competitions and coding, can they reason effectively through complex, open-ended challenges found in frontier physics research? And crucially, what kinds of reasoning tasks do physicists actually want LLMs to assist with? To address these questions, we present the **CritPt** (*Complex Research using Integrated Thinking - Physics Test*, pronounced “critical point”), the first benchmark designed to test LLMs on unpublished, research-level reasoning tasks that broadly covers modern physics research areas, including condensed matter, quantum physics, atomic, molecular & optical physics, astrophysics, high energy physics, mathematical physics, statistical physics, nuclear physics, nonlinear dynamics, fluid dynamics and biophysics. **CritPt** consists of 71 composite research challenges designed to simulate full-scale research projects at the entry level, which are also decomposed to 190 simpler checkpoint tasks for more fine-grained insights. All problems are newly created by over 50 active physics researchers based on their own research. Every problem is hand-curated to admit a guess-resistant and machine-verifiable answer and is evaluated by an automated grading pipeline heavily customized for advanced physics-specific output formats. We find that while current state-of-the-art LLMs show early promise on isolated checkpoints, they remain far from being able to reliably solve full research-scale challenges: the best average accuracy among base models is only 5.7%, achieved by GPT-5 (high), moderately rising to around 10% when equipped with coding tools. Through the realistic yet standardized evaluation offered by **CritPt**, we highlight a large disconnect between current model capabilities and realistic physics research demands, offering a foundation to guide the development of scientifically grounded AI tools.

1 Introduction

Modern physics research encounters increasingly complex systems, specialized tools, and interdisciplinary collaborations [1, 2]. Yet the core standards cannot be compromised: mathematical rigor, true creativity, precise execution and consistency between theory and experiment are all essential. The reality and demanding nature of the subject together set a high entry bar to become a physics researcher and make true breakthroughs more and more difficult to achieve.

Large language models (LLMs) [3–5] show promise in assisting research workflows, for example by identifying relevant literature [6–8], synthesizing scientific knowledge across domains [9–13]. However, these applications largely involve *recombining existing information* and differ fundamentally from the kind of *original reasoning* required to solve research problems in physics. And unlike in natural language tasks where redundancy may mask shallow errors, math and science problems can be unforgiving: a single flawed inference can invalidate the entire solution.

Recently, reasoning-oriented LLMs¹ have made progress on structured multi-step problem solving [21, 17]. These systems usually use encapsulated think tokens as an intermediate process before generating a final answer [17, 22, 23]. They are typically fine-tuned on STEM-focused corpora and optimized for multi-step reasoning tasks, using techniques such as Chain-of-Thought prompting [24], reinforcement learning from verifiable rewards [17], tool use [25] including code execution [26, 27] and web search [28], and scaling inference-time computation [29, 30]. Empirically, these models exhibit behaviors that resemble human reasoning, such as decomposing long problems into coherent substeps, exploring alternatives through trial-and-error, and verifying intermediate results with internal heuristics or external checkers [31, 32, 17]. As a result, these models see striking gains in relatively well-structured reasoning tasks, such as general coding [33–35], high-school academic competitions [36–41], as well as assignments from high school to graduate-level courses [42–47]. However, their performance drops sharply on research-inspired benchmark problems [48–51], where problems are broader in scope and solution spaces are more sparse. So realistically, can the latest LLMs meaningfully assist physicists with the reasoning tasks in frontier research? If so, to what extent?

In this paper, we introduce **CritPt** (*Complex Research using Integrated Thinking - Physics Test*, pronounced “critical point”), a benchmark designed to evaluate LLMs’ reasoning ability in realistic physics research workflows across diverse frontier topics. Our main evaluations are guided by the following lines of inquiry:

- **Can LLMs solve unseen physics research problems beyond their training data?**
The goal of research is to make new discoveries, not to repeat known exercises. While grand open problems are out of reach and hard to verify, can LLMs solve unseen entry-level research problems, where the basic methods and concepts are established, but require nontrivial synthesis and original reasoning to reach a full solution?
- **What reasoning tasks can LLMs help with in realistic physics research workflows today?**
A full-scale research project can often be decomposed into smaller steps or first addressed in a simplified form. In collaborative settings, these modular tasks can be distributed across team members. What types of modular reasoning tasks may LLMs start to assist today?
- **Can we trust LLMs’ reasoning traces and responses in physics research contexts?**
Physics concepts and methods are deeply tied to context-dependent assumptions, where subtle errors in seemingly plausible answers can mislead, especially for those without expert judgment. As a prerequisite check, how reliable are LLMs when tackling complex, unstructured problems in advanced physics, particularly those lying at the boundary of their current capabilities?

CritPt provides a powerful framework for assessing the value of LLMs in realistic physics research workflows, an essential but underexplored component in defining AI’s future role in scientific discovery [52]. In the first release, *CritPt* contains 71 complex, composite *challenges* to simulate full-scale research projects at the entry level, and 190 modular *checkpoints* decomposed from the full challenges to offer more traceable and fine-grained insights on simpler subtasks.

However, designing a benchmark to meet these goals comes with significant practical and technical obstacles. In *CritPt*, we come up with the following design features to address these key obstacles:

¹In this paper, reasoning-oriented models including GPT-5 (high), o3, o4-mini, Gemini 2.5 Pro/Flash, DeepSeek R1 and Claude Opus 4 are evaluated [14–18]. General-purpose chat models, including GPT-5 (minimal), GPT-4o and Llama-4 Maverick [14, 19, 20] are also included for comparison.

- **Frontier research problems standardized by physics experts.**

Research-level problems are underrepresented in LLM benchmarks, as they demand significantly more domain expertise to adapt and validate than textbook-style problems. Consequently, our understanding of AI’s ability in physics tends to center around well-structured problems or focus on a specific discipline, overlooking the complex and open-ended reasoning ability needed for real scientific discovery.

To better represent the depth and breadth of modern physics research, *CritPt* is developed through a 7-month (and ongoing) close collaboration between AI researchers and physics experts from nearly all major physics subfields (Sec. 2.1). Together, we iteratively co-design the content choice, dataset structure and evaluation infrastructure through multiple review stages (Sec. 2.3). All the problems are based on experts’ own research expertise to faithfully represent the **realistic** reasoning demands at the frontier of modern physics, while offering practical signals and accessible insights for the LLM developers.

- **Leakage-resistant and reasoning-focused design.**

Benchmarks sourced directly from public materials or generated from LLMs are susceptible to contamination since such materials are often included in model training data. This potentially leads to inflated performance via memorization or retrieval rather than *genuine reasoning* [53–55], and often suffers from quick saturation and lack of utility value [56]. Further, simple problem formats such as multiple-choice flattens problem complexity, allowing shortcut guessing and easy hacking [57, 58]. Fully open sourced benchmarks can be compromised over time, due to misuse or unintended contamination [59–61].

We mitigate these risks through strict design criteria (Sec 2.2). All problems in the *CritPt* are **unpublished**, hand-curated by physics experts to be well-defined and self-contained with **search-proof** answers. In addition to an **open-ended** question format, problems are constructed to have **guess-resistant** final answers such as arrays of floating-point numbers and complicated symbolic expressions [49]. We publish one example challenge with checkpoints, solutions, and error analysis (Sec. 2.4). The solutions to the other 70 challenges (test set) are kept private to avoid potential contamination.

- **Physics-informed scalable auto-grading pipeline.**

Grading physics problems is traditionally resource-intensive, requiring experts to verify all steps, recognize valid alternative paths, and detect subtle loopholes. Some use LLM judges, which can be unreliable due to sensitivity to superficial factors such as prompt wording or answer format, especially when evaluating content beyond the judge’s own capacity [62, 63]. An alternative is grading the final answer only. However, for open-ended problems in advanced physics, even automating final-answer grading is complicated by technical challenges like parsing free-form LLM outputs and standardizing advanced physics notations [64].

Our evaluation framework is canonical, scalable and physics-informed (Sec. 3). All final answers are **machine-verifiable** by careful design. The LLM answers are first normalized into **structured** code blocks, then scored using **custom scripts** that supports numerical values, SymPy-compatible symbolic expressions [65], and executable Python functions with test cases [48, 50]. Our autograder also accounts for physically meaningful **error tolerances** and equivalent forms specified by physics experts. Though labor-intensive to build upfront, this pipeline enables scalable, high-fidelity evaluation for diverse, complex output formats in advanced physics.

Overall, our physics experts consider *CritPt* challenges comparable in difficulty to the kind of warm-up research exercises that a hands-on principal investigator might assign to junior graduate students, which require solid physics training and some domain expertise, yet remain accessible through thoughtful exploration. In this sense, this benchmark intends to probe the **critical point** of AI reasoning: the transition from producing plausible responses based on superficial pattern recognition, to genuinely reasoning through real-world problems in frontier physics research.

In Sec. 4, we show that current state-of-the-art LLMs are making early progress on isolated checkpoints, but remain far from being able to reliably solve full research-scale challenges. Even the best-performing base model on *CritPt*, GPT-5 (high), reaches only 5.7% average accuracy on challenges, with most other models scoring near zero. When equipped with tools (code interpreter and web search), GPT-5 (high) improves modestly to 12.6% accuracy. More stringent evaluation metrics reveal that current LLMs still lack reliability when tackling research-level problems, underscoring the gap between today’s models and the demands of realistic physics research workflow.

2 Design choices of *CritPt*

We begin by describing the data sources and coverage of *CritPt* in Sec. 2.1, followed by the technical problem criteria in Sec. 2.2. The data creation and review process are outlined in Sec. 2.3, and Sec. 2.4 presents the structure of a *CritPt* challenge with an illustrative example.

2.1 Source and coverage: hand-curated research challenges from the physics community

We source our benchmark data from the problems and reasoning tasks that physicists encounter in real research practice. Because modern physics is highly specialized, this is only possible through a large-scale collaboration with more than 50 physics researchers across 30 institutions worldwide, including senior Ph.D. students, postdocs, and professors. Each contributor crafts problems based on their own research expertise, producing a dataset that has both depth of realistic research and diversity in *disciplines*, *topics* and *flavors* in the modern physics landscape.

Research Area	Challenges	% of Total	Checkpoints	% of Total
Condensed Matter Physics	25	35.2%	70	36.8%
Quantum Information, Science & Technology	17	23.9%	43	22.6%
Atomic, Molecular & Optical	14	19.7%	42	22.1%
High Energy Physics	10	14.1%	30	15.8%
Mathematical Physics	10	14.1%	22	11.6%
Gravitation, Cosmology & Astrophysics	9	12.7%	25	13.2%
Statistical Physics & Thermodynamics	9	12.7%	24	12.6%
Nuclear Physics	7	9.9%	19	10.0%
Nonlinear Dynamics	4	5.6%	12	6.3%
Fluid Dynamics	2	2.8%	6	3.2%
Biophysics	2	2.8%	4	2.1%
Total	71		190	
Covering Multiple Areas	33	46.5%	88	46.3%

Table 1: The physics research *areas* covered by *CritPt*’s challenges and checkpoints.




Figure 1: *CritPt*’s challenges (**left**) and checkpoints (**right**) cover three flavors of physics research – theoretical, experimental, and computational – encountered by physics researchers.

As shown in Table 1, 71 challenges and 190 checkpoints in *CritPt* cover a broad range of modern physics research *disciplines*.² Among them, 33 challenges and 88 checkpoints cover two or more disciplines, reflecting the growing interdisciplinarity of today’s physics research [2]. Within these

²Our categorization is based on a modified version of the Physics Subject Heading (PhySH) classification scheme created by the American Physical Society [66].

disciplines, we cover *topics* ranging from quantum error correction (relevant to industry) to string theory (quest for fundamental particles), from cell dynamics (small scales) to black holes (large scales), from nonlinear optics (experimental techniques) to delicate asymptotics of special functions (math tricks). Most problems are related to experts’ own publications on high-profile physics journals, such as Nature, Science, Physical Review series. For detailed coverage, see the list of challenges in A.1.

CritPt also covers three major *flavors* of physics research: theoretical, experimental, and computational, as shown in Fig. 1. This three-way grouping reflects how physicists commonly describe themselves.³ Here, we cannot exhaustively cover all types of tasks in each category and instead sample a representative cross-section of reasoning tasks. For example, an experimental challenge cannot ask an LLM to run equipment directly but can focus on designing or interpreting an experiment under realistic constraints.

Notably, not all problems in *CritPt* are about polished research questions that one finds in publications. Some problems are inspired by less celebrated but essential aspects of real research, such as failed trials, tedious intermediate calculations, or subtle insights that are rarely documented in papers. These “insider” elements can only be provided by domain experts, and help further differentiate between pattern matching from genuine reasoning. An LLM capable of reliably solving these challenges would mark a major breakthrough in AI for science.

2.2 Benchmark criteria: leakage-resistant and reasoning-focused design

With the source and domain coverage of our data established, the next obstacle is to standardize inherently unstructured research-level problems into a benchmark format that provides accurate signals of genuine reasoning, while allowing scalable evaluation. To guide this process, we define the following technical criteria for constructing *CritPt* problems:

- **Search-proof but solvable.** All problems in *CritPt* are newly created and carefully constructed in a way such that their final answers cannot be retrieved through web search. Meanwhile, they are possible to solve with the publicly known knowledge (i.e., no confidential or private information is needed). All questions are crafted to be well-posed with unambiguous constraints and verifiable final answers. Solving them should demonstrate a deep understanding of the physical scenario, correct application of methods under coherent assumptions, and precise multi-step reasoning and execution. *CritPt*’s problems mainly fall into three categories or their combinations: (1) modified versions of published results to test out-of-distribution generalization, which simulates the realistic research scenario of a follow-up project based on existing results; (2) a non-trivial application of a method to a specific system, which tests understanding and utilizing a method under different physical constraints; (3) non-trivial intermediate steps of a calculation not explicitly shown in a paper, which tests the ability to reproduce published results and understand enough to fill in gaps in the context. We note that these niche contents are also unlikely to appear in future publications, but mirrors frequently occurring tasks in daily research activity.
- **Open-ended Q&A format with verifiable answers.** We adopt open-ended question formats with various answer formats, mostly numbers or symbolic expressions. All the symbols, conventions and physical units are explicitly given in the problems to prepare for canonical grading later. If an answer expression is too complicated for reliable symbolic manipulation in SymPy or admits too many equivalent forms, we ask models to return a Python function as the answer and evaluate it with test cases. In rare cases that asking a question with a binary or categorical answer (e.g., “Yes/No”) is essential, we ask a set of related questions and consider the model’s solution correct only when all are answered correctly, mirroring how real scientific understanding often requires consistency across multiple angles.
- **Guess-resistant construction tailored for physics contents.** Though search-proof by construction, physics results, particularly the elegant and memorable ones, often take on some commonly occurring values, such as 0, $1/2$ or π , regardless of the system variation or the derivation path. For example, in condensed matter physics, many systems are extremely complex, but universal quantities such as topological numbers are often shared by systems with different microscopic details. To mitigate the risk of guessing, we carefully choose the physical systems and the quantities to ask that distinguish between correct and incorrect reasoning paths, to ensure that models must follow the intended sequence of physical reasoning to arrive at the correct conclusion. Each final answer usually contains

³This also aligns with the *technique* facet of the PhySH classification [66].

at least one non-universal quantity in a complicated format, such as floating-point numbers with several-decimal precision, large integers or dimension-dependent symbolic expressions.

We note that our design criteria on answer formats are partially inspired by SciCode [48], FrontierMath [49] and TPBench [50].

2.3 Quality control: iterative development and multi-level expert review

Guided by the benchmark criteria above, each challenge in *CritPt* goes through an extensive iterative creation and a multi-stage review process. Every data contributor, benchmark coordinator, problem reviewer and scientific writer, holds a Ph.D. in physics or is an active physics Ph.D. student engaged in frontier physics research.

The data collection follows the workflow below:

1. **Initial creation:** The coordinators first provide each physics expert annotator with an at least hour-long introduction to LLMs and the benchmark design criteria. Experts then create problems based on their research expertise. Each submission includes a solution often more detailed than a typical journal paper, containing step-by-step explanations, algebraic derivations, numerical codes, supporting data, references, and occasionally alternative solutions.
2. **Iterative revision:** The initial draft of each problem undergoes an iterative reviewing process between the expert and coordinators, typically with three or more rounds and up to ten for extremely complex cases. AI researchers and physics experts also jointly analyze LLM responses to make sure authentic, domain-relevant reasoning is being tested rather than spurious artifacts, such as formatting issues, ambiguous prompting, or subtle loopholes. We avoid cherry-picking based on specific behavior of a particular model to ensure fair comparisons and long-term utility.
3. **Expert review:** After iterative reviewing, each problem undergoes high-level peer review by researchers in closely related areas, while technical derivations and algebraic steps are validated by additional physics experts. Final write-ups are edited by a science writing specialist for clarity and accessibility.

On average, it takes more than 40 hours of expert effort to create one full challenge in *CritPt*. All *CritPt*'s contributors and consultants (see Acknowledgment) are given access to leading LLMs and encouraged to experiment with them. Experts' first-hand observations of model performance, limitations, and behaviors have deeply transformed our benchmark design throughout a 7-month collaboration. As a result, *CritPt* not only reflects realistic reasoning demands that physicists themselves care about, but is also an effort to provide direct and actionable feedback for AI developers, including those without an advanced physics background.

2.4 Structure of a challenge: an example

We illustrate the structure and design of a full *CritPt* challenge with an example, "Quantum Error Detection", in this section.

Each **challenge** is designed as a self-contained, research-style problem at the level of a junior researcher. The *setup* section provides all necessary background and context, mimicking how a mentor would define the scope and clarify assumptions when on-boarding the researcher. The *challenge question* then poses the central research inquiry within this setup, requiring complex multi-step original reasoning to solve.

Although most contributing experts agree that current LLMs lag far behind human-level research reasoning in their fields, they also recognize the potential for models to assist with more focused or granular tasks. To capture this, each *CritPt* challenge is decomposed into 2–4 **checkpoint** questions to isolate specific reasoning steps. While this decomposition naturally reduces complexity, these checkpoints still preserve the depth and reasoning demands of realistic research workflow, going well beyond mechanical steps like formula substitutions. They include but are not limited to: filling in intermediate steps in a long derivation, solving simplified versions of the full task (e.g., 1D case before higher dimensions), or analyzing relevant special cases (e.g., behaviour in the high-temperature limit). If LLMs can reliably handle such tasks, they can save physicists significant effort and speed up scientific discovery. We note that the *setup* section is also provided to the models when evaluating checkpoints.

Challenge: Quantum Error Detection

Setup:

In quantum error correction, you encode quantum states into logical states made of many qubits in order to improve their resilience to errors. In quantum error detection, you do the same but can only detect the presence of errors and not correct them. In this problem, we will consider a single $[[4,2,2]]$ quantum error detection code, which encodes two logical qubits into four physical qubits, and investigate how robust logical quantum operations in this code are to quantum errors.

Our convention is that the four physical qubits in the $[[4,2,2]]$ code are labelled 0,1,2,3. The two logical qubits are labelled A and B. The stabilizers are $XXXX$ and $ZZZZ$, where X and Z are Pauli matrices. The logical X and Z operators on the two qubits are $X_A = XIXI$, $X_B = XXII$, $Z_A = ZZII$, $Z_B = ZIZI$, up to multiplication by stabilizers.

We will consider different state preparation circuits consisting of controlled not $CNOT_{ij}$ gates, where $CNOT_{ij}$ has control qubit i and target qubit j . As a simple model of quantum errors in hardware, we will suppose that each $CNOT_{ij}$ gate in the circuit has a two qubit depolarizing error channel following it that produces one of the 15 non-identity two-qubit Paulis with equal probability $p/15$. The probability p indicates the probability of an error in a single two-qubit gate. We will assess the logical infidelity of certain state preparation protocols as a function of the physical infidelity p .

Challenge question:

Suppose that we prepare a logical two-qubit $|00\rangle_{AB}$ state in the $[[4,2,2]]$ code. To do so, we introduce an ancilla qubit, qubit 4, and use the following state preparation circuit:

$$M_4(CNOT_{04})(CNOT_{34})(CNOT_{23})(CNOT_{10})(CNOT_{12})(H_1)$$

Note that this equation is written in matrix multiplication order, while the quantum operations in the circuit occur in the reverse order (from right-to-left in the above equation). H is a single-qubit Hadamard gate and M is a single-qubit measurement. The ancilla is used to detect errors in the state preparation circuit and makes the circuit fault-tolerant. If the ancilla measurement is $|0\rangle$ ($|1\rangle$), the state preparation succeeds (fails).

What is the logical state fidelity of the final 2-qubit logical state at the end of the circuit as a function of two-qubit gate error rate p , assuming the state is post-selected on all detectable errors in the code and on the ancilla qubit measuring $|0\rangle$?

Answer:

$$F_{\text{logical}} = 1 - \frac{\frac{16}{25}p^2 - \frac{128}{125}p^3 + \frac{2048}{3375}p^4 - \frac{32768}{253125}p^5}{1 - \frac{68}{15}p + \frac{704}{75}p^2 - \frac{32768}{3375}p^3 + \frac{253952}{50625}p^4 - \frac{262144}{253125}p^5}$$

Checkpoints

Checkpoint 1:

Suppose that we wish to prepare a logical two-qubit GHZ state $(|00\rangle_{AB} + |11\rangle_{AB})/\sqrt{2}$ in the $[[4,2,2]]$ code. To do so, we use the following state preparation circuit:

$$(CNOT_{03})(H_0)(CNOT_{21})(H_2).$$

Note that this equation is written in matrix multiplication order, while the quantum operations in the circuit occur in the reverse order (from right-to-left in the above equation). H is a single-qubit Hadamard gate.

What is the physical state fidelity of the final physical 4-qubit state at the end of the circuit as a function of the two-qubit gate error rate p ?

Answer:

$$F_{\text{physical}} = \left(1 - \frac{12}{15}p\right)^2$$

Checkpoint 2:

Suppose that we wish to prepare a logical two-qubit GHZ state $(|00\rangle_{AB} + |11\rangle_{AB})/\sqrt{2}$ in the $[[4,2,2]]$ code. To do so, we use the following state preparation circuit:

$$(CNOT_{03})(H_0)(CNOT_{21})(H_2).$$

Note that this equation is written in matrix multiplication order, while the quantum operations in the circuit occur in the reverse order (from right-to-left in the above equation). H is a single-qubit Hadamard gate.

What is the logical state fidelity of the final 2-qubit logical state at the end of the circuit as a function of the two-qubit gate error rate p , assuming the state is post-selected on all detectable errors in the code?

Answer:

$$F_{\text{logical}} = 1 - \frac{\frac{16}{75}p^2}{1 - \frac{8}{5}p + \frac{64}{75}p^2}$$

Checkpoint 3:

Suppose that we prepare a logical two-qubit $|00\rangle_{AB}$ state in the $[[4,2,2]]$ code. To do so, we introduce an ancilla qubit, qubit 4, and use the following state preparation circuit:

$$M_4(CNOT_{04})(CNOT_{34})(CNOT_{23})(CNOT_{10})(CNOT_{12})(H_1)$$

Note that this equation is written in matrix multiplication order, while the quantum operations in the circuit occur in the reverse order (from right-to-left in the above equation). H is a single-qubit Hadamard gate and M is a single-qubit measurement. The ancilla is used to detect errors in the state preparation circuit and makes the circuit fault-tolerant. If the ancilla measurement is $|0\rangle$ ($|1\rangle$), the state preparation succeeds (fails).

What is the logical state fidelity of the final 2-qubit logical state at the end of the circuit as a function of two-qubit gate error rate p , assuming the state is post-selected on all detectable errors in the code and on the ancilla qubit measuring $|0\rangle$?

Answer:

$$F_{\text{logical}} = 1 - \frac{\frac{16}{25}p^2 - \frac{128}{125}p^3 + \frac{2048}{3375}p^4 - \frac{32768}{253125}p^5}{1 - \frac{68}{15}p + \frac{704}{75}p^2 - \frac{32768}{3375}p^3 + \frac{253952}{50625}p^4 - \frac{262144}{253125}p^5}$$

See design ideas behind the example challenge in [A.5.1](#) and detailed solutions on this [web page](#).

3 Evaluation pipeline

We implement a high-fidelity, automated evaluation framework combining a structured two-step generation protocol (Sec. 3.1) with a robust, canonical grading system (Sec. 3.2). This setup ensures both faithful assessment of reasoning quality and rigorous, scalable verification across diverse output formats.

3.1 Two-step answer generation from models

To disentangle the reasoning process from answer formatting, we adopt a two-step generation strategy, sketched in Fig. 2 (Left):

- In the first step, the model is prompted to generate a complete solution using free-form natural language and mathematical derivations (see [A.2](#) for the system prompt). This allows the model to reason without constraints imposed by output formatting templates.
- In the second step, we guide the model to extract and standardize its final answer into a designated code block template we provide (see [A.2](#) for the parsing prompt and the example template). This template enforces a canonical structure suitable for parsing and grading.

By separating solution reasoning from formatting, we avoid premature conversions (e.g., via SymPy) that may distort intermediate steps, and reduce parsing errors from inconsistent model output styles.




Figure 2: A schematic overview of the two-step generation process and the grading system. **Left:** The two-step generation protocol separates problem-solving (first round) from answer formatting (second round). **Right:** The automated grading system compares the model output against the gold answer from experts using scripts customized according to the expected answer format.

3.2 Auto-grading system

We extract the final answer by parsing the Python code block generated in the second step, which contains an answer function. Parsing is designed to tolerate surrounding text while enforcing strict formatting within the code block. To ensure safety and compatibility, only a limited set of libraries (e.g., `math`, `numpy`, `sympy`, `scipy`) is allowed, and potentially unsafe operations are filtered.

Next, we compare the extracted model’s answer function with the gold answer provided by physics expert through an automated grading system. Our automated grading system supports three primary answer types: numerical values, symbolic expressions, and Python functions, each with their own evaluation logic, sketched in Fig. 2 (Right):

- **Numerical values:** We assess correctness against gold answers using expert-provided tolerance ranges that generously account for the physical or numerical sensitivity of each problem. For exact results, a default precision of 12 significant digits is enforced to accommodate potential floating-point errors introduced by conversion to Python code.
- **Symbolic expressions:** For problems involving symbolic algebra, we use a hierarchical grading script based on SymPy, implemented underneath the `sympy_grading()` function in Fig. 2. It starts with built-in equivalence checking and algebraic simplification, extended with custom routines tailored to the structure of each expression when standard simplification proves insufficient. We also accommodate issues that can be easily overlooked such as math object type conflicts in SymPy.
- **Python codes:** When models return executable functions (e.g., for parametric solutions), we grade the model answer using curated test cases selected by physics experts [48]. These cases are chosen to probe physically meaningful parameter ranges and edge conditions.

For composite answers (e.g., tuples or dictionaries of results), we apply element-wise grading, and an answer is considered correct only if all components match the expert answer.

To ensure runtime safety and isolation, each answer is executed in a sand-boxed environment with enforced resource limits: 30-second wall-clock timeout and bounded memory usage (or the resource limits specified by the expert for computationally heavy tasks; whichever is higher between the default and expert’s specification). This prevents pathological behaviors such as infinite loops or excessive allocation created by models, while still supporting computation-intensive problems.

4 Results

In this section, we evaluate 10 state-of-art models (configuration in Table 2) to directly answer the three lines of inquiry in the introduction. Evaluation is performed *independently* on two levels: the full

challenges (Sec. 4.1) and finer-grained **checkpoints** (Sec. 4.2). In these two sections, we report **average accuracy** over five runs as our primary metric to reduce high stochasticity in model behavior. Next in Sec. 4.3, we adopt a stricter metric, **consistently solved rate**, to probe reliability of model performance. Table 3 is a brief summary of these aggregated accuracy statistics. Detailed analysis of model responses beyond aggregated statistics is discussed in Sec. 4.4.

Model Name	Reasoning Effort	Tool Use	Company
GPT-5 (high)	high	/	OpenAI
GPT-5 (high, code)	high	code interpreter	OpenAI
GPT-5 (high, code & web)	high	code interpreter, web search	OpenAI
o3 (high)	high	/	OpenAI
o4-mini (high)	high	/	OpenAI
Gemini 2.5 Pro	reasoning_tokens=27000	/	Google
Gemini 2.5 Flash	Default	/	Google
DeepSeek R1	Default	/	DeepSeek
Claude Opus 4	reasoning_tokens=27000	/	Anthropic
GPT-5 (minimal)	minimal*	/	OpenAI
Llama-4 Maverick	/	/	Meta
GPT-4o	/	/	OpenAI

* GPT-5 (minimal) sets reasoning_effort = minimal, which means zero reasoning tokens used.

Table 2: Models and their API configurations used in our evaluation. Both reasoning-oriented models (white background) and general-purpose chat models (grey background) are included. More details in A.3.

Model Name	Average Accuracy (%)			Consistently Solved Rate (%)		
	Challenge	Checkpoint		Challenge	Checkpoint	
		w/o expert answer	w/ expert answer		w/o expert answer	w/ expert answer
GPT-5 (high, code & web)	12.6	21.4	24.5	10.0	14.4	17.6
GPT-5 (high, code)	10.6	20.0	23.9	8.6	15.5	17.6
GPT-5 (high)	5.7	15.3	20.0	4.3	10.7	15.5
Gemini-2.5 Pro	2.0	8.1	9.1	0.0	4.3	3.7
o3 (high)	1.4	7.4	10.6	0.0	3.7	5.9
DeepSeek R1	1.1	5.1	6.6	0.0	2.1	3.2
Gemini-2.5 Flash	1.1	4.1	4.6	0.0	2.1	2.1
o4-mini (high)	0.6	5.6	7.1	0.0	2.7	2.1
Claude Opus 4	0.3	2.6	3.0	0.0	0.5	1.6
GPT-5 (minimal)	0.0	2.6	3.7	0.0	1.6	2.1
Llama-4 Maverick	0.0	2.2	2.1	0.0	1.1	0.5
GPT-4o	0.0	1.7	1.2	0.0	0.0	0.0

Table 3: A high-level summary of *CritPt* evaluation results.

4.1 Challenge-level evaluation: can LLMs solve unseen research problems?

We first assess model performance on *CritPt* challenges without intermediate supervision, testing their end-to-end complex reasoning in unseen full-scale research problems. As shown in Fig. 3, all models score low in terms of average accuracy of five independent runs across 70 challenges in *CritPt* test set. Among base models (no external tools), GPT-5 (high), achieves only 5.7% as the best result, with all other models below or equal to 2%.




Figure 3: A comparison of 10 models’ performance on 70 test *CritPt* challenges. Each model is tested on every challenge in five independent runs. **Main plot:** the average accuracy over all runs and all challenges for each model. **Inset a:** the average number of reasoning tokens used per run for each model. **Inset b:** the average cost (USD) per run, calculated from token usage and API pricing for each model (A.4).

Tool use can offer modest but meaningful gains. With access to the code interpreter, GPT-5 (high, code) improves to 10.6%, a notable relative improvement from the base model, consistent with the importance of numerical tools in modern physics research. Adding web search brings only a marginal increase to 12.6%, confirming that *CritPt*’s search-proof design effectively resists short-cutting via retrieval and emphasizes genuine reasoning.

As expected, all the general-purpose chat models score zero, while reasoning-oriented models start making progress, however small, suggesting some emerging capability from explicitly structured reasoning processes. Meanwhile, with long reasoning chains and more verbose outputs, these reasoning-oriented models consume significantly more resources. As shown in **Inset a** of Fig. 3, these models consume much more tokens per run than general-purpose models. Note that GPT-5 (high, code & web) uses an order of magnitude more tokens than other reasoning models, primarily caused by large amounts of web-retrieved contents counted as extra input tokens, not necessarily indicating deeper thinking. For a more detailed token consumption breakdown by type and number, please see A.4. Overall, this shows that all reasoning-oriented models engage deeply with these challenges, which are indeed difficult even for LLMs with long context windows and extensive computational resources.

Inference cost can be another critical constraint for scaling LLM usage in research. **Inset b** of Fig. 3 reports the average cost per run on *CritPt* challenges, determined by total token usage and each vendor’s standard API pricing (see A.4 for details). The large token consumption by advanced models naturally drives up cost, but the gain in performance is often disproportionate. This highlights the need for *efficient* reasoning: performance gains should not rely solely on extended context length, especially in scientific domains where the solution space is sparse and brute-force exploration is ineffective. Pricing policies further impact cost, as high-performing models are typically commercial. For example, DeepSeek R1 is relatively affordable thanks to its low per-token rates, whereas Claude Opus 4 is the most expensive due to its premium pricing.

In summary, current LLMs are far away from being able to solve unseen physics research problem at a junior-researcher level. While coding tools help make a small initial step, they are not sufficient to overcome fundamental reasoning bottlenecks. Notably, GPT-5 was released near the very end of our data collection cycle, but only shows limited progress over its predecessor, o3 (high), on *CritPt* challenges. This suggests that a realistic benchmark like *CritPt* has the potential to resist new generations of models for a substantial amount of time.

4.2 Checkpoint-level evaluation: smaller tasks that LLMs can assist today?

To better differentiate current model capabilities and isolate failure modes, each *CritPt* challenge is decomposed into 2–4 checkpoint questions. These checkpoints are evaluated sequentially in a multi-turn conversation format, simulating how researchers might naturally interact with an assistant system during problem solving.

We use two evaluation setups for checkpoints, illustrated in Fig. 4, to test reasoning effectiveness in different research scenarios:

- **Self-carryover (without expert answer):** The model proceeds sequentially, using only its own previous outputs (Fig. 4a). This setting captures a realistic scenario where the overall problem is decomposable, but intermediate results remain uncertain and errors can propagate and compound.
- **Oracle carryover (with expert answers):** The model also proceeds sequentially, but is given ground-truth answer to the prior checkpoint before attempting the next (Fig. 4b). This setup intends to test isolated local tasks by removing upstream error effects, or assess whether a model can effectively use answers to relevant tasks as hints.




Figure 4: Schematic of the two experimental setups for evaluating sequential checkpoints within a multi-turn conversation. (a) Self-carryover without expert answer: The model’s own output from the previous checkpoint is used as context for the next one. (b) Oracle carryover with expert answers: The correct answer (shown in red) to the previous checkpoint is provided before the model attempts the next checkpoint.

As shown in Fig. 5, current LLMs show early promise in assisting with checkpoints, more localized or well-scoped tasks compared to challenges. GPT-5 (high) again outperforms all other models. In the self-carryover without expert answer (solid bar), GPT-5 (high) reaches 15.3%. This improves to 20.0% with code interpreter and reaches 21.4% when also equipped with web search. The next best performers are Gemini 2.5 Pro (8.1%) and o3 (high) (7.4%), followed by o4-mini (high) (5.6%), DeepSeek R1 (5.1%) and Gemini 2.5 Flash (4.1%). The remaining models fall below 3%, where small absolute differences may reflect statistical noise rather than meaningful performance gaps, so the ordering should be interpreted with caution.

Almost all models benefit from expert answer injection to earlier checkpoints. The largest gains are observed for GPT-5 (high) family and o3 (high), which improve by more than 3% in the oracle-carryover setting (hatched bar), indicating the ability to leverage correct intermediate results to improve downstream reasoning.




Figure 5: A comparison of 10 models’ performance on the 187 test *CritPt* checkpoints. The average accuracy is aggregated over all runs and all checkpoints, in two setups respectively. Solid bar reports the self-carryover without expert answer, while the hatched pattern reports oracle carryover with expert answers.

Though still difficult, *CritPt* checkpoints seems to fall within the improving front for next generation leading models. This aligns with the feedback from our physics experts, particularly theorists, who have begun cautiously experimenting with LLMs in their daily research, and occasionally find them correct for small and well-defined reasoning tasks. However, LLMs are not fully correct most of the time, so experts must dedicate considerable time to verifying the convoluted yet plausible outputs. This process can sometimes exceed the time required to solve the problem independently. These observations motivate the next section, an attempt to assess reliability of LLM performance with a stricter evaluation metric.

4.3 Reliability metric: can we trust LLM outputs?

In the regime of highly complex and open-ended problems, a robust and trustworthy reasoning process is particularly important. Subtle mistakes from hallucination can be hard to identify and can mislead users who are learning new things and lack expert-level judgment. As a prerequisite check of the models’ reliability, we introduce a stricter performance metric: a problem is deemed as **consistently solved** only if at least four out of five runs give correct final answers.⁴

Applying this criterion leads to a sharp drop in performance across all models, suggesting high stochasticity in model behavior in complex physics research context.

At the full challenge level (Fig. 6a), only GPT-5 (high) is able to solve any problems consistently [67]. The base model scores only 4.3% (3 out of 70 challenges). With tool use, this improves to 8.6% (code) and 10.0% (code & web). All other reasoning-oriented models drop to zero under this stricter measure.

At the checkpoint level (Fig. 6b), leading models are able to consistently solve a very limited number. Interestingly, GPT-5 (high, code) outperforms GPT-5 (high, code & web) with a small margin, further supporting the search-proof design of *CritPt* even at the checkpoint level.

These results imply that current LLMs cannot yet be trusted to reason consistently in high-stakes research contexts. Improving reasoning reliability through better uncertainty calibration, more powerful external validations or other advances remains an open challenge. Meanwhile, the message to the physics community is clear: while LLMs may be useful for exploring or prototyping small subtasks, caution and expert oversight may be necessary in advanced research contexts.

⁴Caveats: this is a heuristic reliability check based on the number of runs we have given our resource constraints. It can also be viewed as a 4/5 super majority vote (or pass@4/5). We are *by no means* claiming statistical sufficiency from five samples nor that 80% accuracy implies high reliability. We welcome sponsorship or third-party host to run more tests.




Figure 6: Percentage of *CritPt* problems consistently solved by models. A problem is considered consistently solved if at least four out of five independent runs yield the correct final answer. (a) Percentage of challenges consistently solved. (b) Percentage of checkpoints consistently solved.

4.4 Detailed analysis of full model responses

Beyond aggregated accuracy metrics, we further analyze model behavior at the level of individual challenges and walk through the full model responses (including reasoning traces when available) with physics experts, which helps surface more qualitative insights not captured by final answer correctness alone.

To support an efficient review process with experts and overcome the limitations of traditional LLM infrastructure, we develop an interactive visualization platform. This tool allows experts to quickly browse model responses, systematically compare performance across tasks and model families, and identify error patterns or interesting behaviors. Table 4 illustrates the visualization of the example challenge, where we can immediately observe an unexpected behavior: adding web search tool on top of coding for GPT-5 (high) actually degrades the model performance for this particular challenge. The full interactive demo including all model responses is available at critpt.com/example.

This interface streamlines non-AI researchers’ engagement with LLM outputs at scale, by abstracting away the technical complexity of API access and evaluation infrastructures, thereby making it easier to provide structured feedback. For the example challenge, the expert’s detailed feedback is provided in A.5.2. The expert notes that although GPT-5 (high) achieves the highest final-answer accuracy, its output formatting is cluttered and difficult to follow, potentially limiting its usefulness in realistic workflows, and this may be a solvable problem in natural language processing. Another interesting observation is that GPT-5 (high) often calls tools in ways that diverge from expert expectations: using code execution even when an analytical solution would be simpler, or performing excessive web retrieval

Model Name	Challenge	Checkpoint 1	Checkpoint 2		Checkpoint 3	
			w/o expert answer	w/ expert answer	w/o expert answer	w/ expert answer
GPT-5 (high, code & web)	2/5	5/5	5/5	5/5	3/5	3/5
GPT-5 (high, code)	4/5	5/5	5/5	5/5	5/5	4/5
GPT-5 (high)	0/5	5/5	5/5	5/5	0/5	0/5
Gemini-2.5 Pro	0/5	3/5	3/5	1/5	0/5	0/5
o3 (high)	0/5	5/5	2/5	2/5	0/5	0/5
DeepSeek R1	0/5	3/5	3/5	2/5	0/5	0/5
o4-mini (high)	0/5	5/5	1/5	0/5	0/5	0/5
Gemini-2.5 Flash	0/5	2/5	0/5	0/5	0/5	0/5
Claude Opus 4	0/5	0/5	0/5	0/5	0/5	0/5
GPT-5 (minimal)	0/5	0/5	0/5	0/5	0/5	0/5
Llama-4 Maverick	0/5	0/5	0/5	0/5	0/5	0/5
GPT-4o	0/5	0/5	0/5	0/5	0/5	0/5

Table 4: Detailed breakdown of evaluation results for the example challenge, “quantum error detection.”

before assessing relevance. These behaviors reveal a gap between current LLM decision heuristics and human-expert research intuition.

A broader analysis of failure modes and emergent reasoning behaviors in different types of physics research problems is ongoing and will be reported in a forthcoming paper.

5 Conclusion

In this work, we introduce *CritPt*, a physics reasoning benchmark designed by experts to probe the capacity of LLMs to meet the authentic reasoning demands of frontier research. Moving beyond traditional evaluation formats, *CritPt*’s self-contained challenges mimic how a mentor frames a problem for a junior researcher, while its modular checkpoints allow for a fine-grained analysis of reasoning capabilities without sacrificing scientific depth. By assembling a diverse set of unpublished, guess-resistant problems developed by 50+ active physicists, we provide the first systematic testbed grounded in the realistic workflows, complexity, and failure sensitivity of modern physics.

Our evaluation results demonstrate a striking gap between current model performance and the depth, rigor, creativity and precision required for physics research. While advanced reasoning models such as GPT-5 (high) show early promise on narrowly scoped tasks, particularly when augmented with coding tools, their ability to maintain coherence and correctness through full-scale challenges remains minimal. Furthermore, our reliability analysis shows that even correct answers are often not consistently reproducible, highlighting a severe deficit in the robustness required for high-stakes scientific applications.

Beyond a benchmark, *CritPt* represents a collaborative bridge between the physics and AI communities. For AI developers, it translates the abstract reasoning demands of physics research into a standardized dataset and provides concrete and accessible signals via a physics-informed automatic evaluation pipeline, providing immediate high-quality feedback for model development. For physicists, it offers a grounded introduction to the current capabilities and limitations of AI assistants, supported by an interactive visualization tool that enables domain experts to efficiently inspect, analyze, and critique model outputs at scale. This process creates a shared vocabulary, clarifying the requirements for future systems that aim not just to answer textbook questions, but to engage with the complex and iterative workflows of scientific discovery.

Looking forward, *CritPt* is a powerful and scientifically grounded framework to measure progress. It offers a clear metric of how far AI has come and, more importantly, how far it has yet to go to augment genuine discovery in physics. We hope this benchmark will not only guide the development of more capable and reliable reasoning models but also catalyze deeper conversations between the physics and AI communities on what it truly means for a machine to reason in a scientific context.

Acknowledgments and Disclosure of Funding

We gratefully acknowledge insightful discussions with Nigel Goldenfeld and Carl M. Bender on physics reasoning and current AI limitations. We also thank Zixiang Lu, Qiuling Fan, Semih Kacmaz, Xiaoning Wang, Sang Hyun Choi, Ethan Lake, Tong Wang, Carlo Siebenschuh, for their help on problem development. We are appreciative of conversations with Xiao Chen, Andrew Lucas, Stephen Taylor, and Helvi Witek regarding problem formulation and student mentorship. Finally, we thank the following physicists for sharing their experiences and perspectives on AI (in alphabetical order): Botao Du, Andrew Guo, Peter Johnsen, Zhiru Liu, Takumi Matsuzawa, Chenyu Tian, Zihan Wang, and Bo Zou.

This work was supported by Laboratory Directed Research and Development (LDRD) funding from Argonne National Laboratory, provided by the Director, Office of Science, of the U.S. Department of Energy (DOE) under Contract No. DE-AC02-06CH11357. We gratefully acknowledge computing resources provided by: the Argonne Leadership Computing Facility (ALCF), a DOE Office of Science User Facility under Contract DE-AC02-06CH11357; the Delta advanced computing and data resources supported, by the National Science Foundation (NSF) (award OAC 2005572) and the State of Illinois. DeltaAI (award OAC 2320345) is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications. Eliu Huerta also acknowledges partial support from NSF grants OAC-2209892 and OAC-2514142. Yong Zhao’s material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics through Contract No. DE-SC0012704, No. DE-AC02-06CH11357, within the framework of Scientific Discovery through Advanced Computing (SciDAC) award Fundamental Nuclear Physics at the Exascale and Beyond, and under the umbrella of the Quark-Gluon Tomography (QGT) Topical Collaboration with Award DE-SC0023646.

A Appendix

A.1 List of *CritPt* challenges

[Example challenge](#) [68–71]:

- Quantum error detection

[Test set \(70 challenges\)](#) [72–315]:

1. Holographic Weyl anomaly
2. Population growth rate from stochastic model of growth and cell-size regulation
3. Geodesic in AdS/BCFT of a black hole
4. Orbital angular momentum conservation in high harmonic generation
5. Marchenko-Pastur entropy
6. Twisted Bilayer MoTe_2
7. Noisy quantum sensing
8. Scalar spectrum in Nieh–Yan gravity
9. Nieh–Yan modified gravity with Torsion in FRW space-time
10. Axion inflation with Chern-Simons
11. Gapped edge of the Moore-Read state
12. Parafermion zero modes tunneling
13. Verlinde lines in the Moore-Read CFT
14. High/low-temperature duality in Ising Torus
15. Decohered Affleck-Kennedy-Lieb-Tasaki (AKLT) model
16. Interacting Chern insulator
17. Zero temperature entropy in Sachdev-Ye-Kitaev models
18. Optical binding force
19. Cascade optical parametric amplifier
20. Torsional levitated optomechanics
21. Numerical LaMET matching
22. Single particle Holevo Information
23. One-loop correction of quasi-PDF
24. LaMET matching in Coulomb gauge
25. Minimum Doppler factor of a relativistic jet
26. Spherical cavity shifts
27. One-axis twisting model with dissipation
28. Optical conductivity of the Hubbard model
29. Hubbard model in an optical lattice
30. Orthogonal non-isometric maps
31. Linear stability analysis of Rayleigh-Bénard convection
32. Rayleigh-Darcy convection with mixed boundary conditions
33. Condensation of three types of particles
34. Quantum tensor networks
35. Quantum inverse problem
36. Oscillation amplitude in transient dynamics of autocatalytic cycles
37. Quantum geometry

38. Scattering rate of the Hatsugai-Kohmoto model
39. Alteration of cavity field coherences due to atom-cavity interaction
40. Hydrodynamic modes in Schwinger-Keldysh
41. Energy in many body quantum systems
42. Graphene minimal conductivity
43. Disclination charge
44. Ground state in Kitaev honeycomb model
45. Goniopolarity in semiconductor
46. PXP scar
47. Integrals of motion
48. Lattice Gaussian sum
49. Long-range light cone
50. Many-body NC Partitions
51. Random walk $a_3(t)$
52. Efimov effect in three body problem
53. Extended obstruction tensors
54. Magic wavelengths for Yb isotopes
55. INS cross-section for scattering from an oscillator
56. Dark matter detection with Cosmic Explorer
57. LIGO with modified mirrors for ultralight vector dark matter
58. Magnetic space group identification
59. Charge density wave diffraction
60. Superresolution
61. Quantum search time
62. Quantum games in multi-slit interference
63. Noise robustness in Kochen–Specker and Spekkens contextuality
64. Conformal correlators
65. Constructing fermionic matrix operators
66. Generating function of index
67. Quantum capacity for quantum channels
68. Quantum f -divergence
69. Contraction coefficients for quantum relative entropy
70. Row-Twirling channel on qubit lattice

A.2 Prompts for two-step answer generation from model

Step 1: system prompt for full problem solving

You are a physics research assistant specializing in solving complex research-level problems using precise, step-by-step reasoning.

Input

Problems will be provided in Markdown format.

Output (Markdown format)

1. **Step-by-Step Derivation** – Show every non-trivial step in the solution. Justify steps using relevant physical laws, theorems, mathematical identities (or numerical codes) ^a.
2. **Mathematical Typesetting** – Use LaTeX for all mathematics: $\$ \dots \$$ for inline expressions, $\$\$ \dots \$\$$ for display equations.
3. **Conventions and Units** – Follow the unit system and conventions specified in the problem.
4. **Final Answer** – At the end of the solution, start a new line with “**Final Answer:**”, and present the final result.
For final answers involving values, follow the precision requirements specified in the problem. If no precision is specified:
 - If an exact value is possible, provide it (e.g., $\sqrt{2}$, $\pi/4$).
 - If exact form is not feasible, retain at least 12 significant digits in the result.
5. **Formatting Compliance** – If the user requests a specific output format (e.g., code, table), provide the final answer accordingly.

^aContent in the parenthesis is only given when code interpreter is enabled.

Step 2: parsing prompt for answer formatting

Populate your final answer into the code template provided below. This step is purely for formatting/display purposes. No additional reasoning or derivation should be performed. Do not import any modules or packages beyond what is provided in the template.

```
1 import sympy as sp
2 p = sp.symbols('p')
3
4 def answer(p):
5     r"""
6     Return the expression of the logical state fidelity
7     in SymPy format.
8
9     Inputs
10    -----
11    p: sympy.Symbol, two-qubit gate error rate, $p$
12
13    Outputs
14    -----
15    F_logical: sympy.Expr, logical state fidelity
16    """
17
18    ----- FILL IN YOUR RESULTS BELOW -----
19    F_logical = ... # a SymPy expression of inputs
20    -----
21
22    return F_logical
```

A.3 Detailed API setup used in evaluation

We evaluate eight proprietary models across four major providers (OpenAI, Google, DeepSeek, Anthropic) using Inspect AI [316], and one open-source model (Meta Llama-4 Maverick) via Together AI [317].

From OpenAI, we evaluate four models. These include their flagship reasoning model, GPT-5 [14], as well as two earlier-generation models from the o-series: the large o3, designed for extended multi-step reasoning, and its lighter variant, o4-mini, optimized for latency and cost [15]. We also assess GPT-4o, a general-purpose multimodal model [19]. From Google, we evaluate Gemini 2.5 Pro, their most advanced “thinking model” to date, which excels at complex reasoning, coding, and scientific tasks [16]. We also include Gemini 2.5 Flash, a lightweight, high-throughput variant engineered for speed and efficiency, while retaining core reasoning capabilities [16]. From DeepSeek, we evaluate DeepSeek R1, an open-weight model trained with multi-stage reinforcement learning to enable multi-step reasoning and self-reflection [17]. It achieves performance comparable to OpenAI-o1-1217 [21]. From Anthropic, we include Claude Opus 4, their flagship model for complex coding and reasoning, with improved precision in instruction-following [18]. Finally, from Meta, we evaluate Llama-4 Maverick, an open-weight, instruction-tuned multimodal model with strong performance on a broad range of widely reported benchmarks [20].

For reproducibility, model configurations including API names and configurations are summarized in Table 5.

Model Name	API Name	Reasoning Effort	Tool Use
GPT-5 (high, code & web)	gpt5-2025-08-07	high	code interpreter, web search
GPT-5 (high, code)	gpt5-2025-08-07	high	code interpreter
GPT-5 (high)	gpt5-2025-08-07	high	/
o3 (high)	o3-2025-04-16	high	/
o4-mini (high)	o4-mini-2025-04-16	high	/
Gemini 2.5 Pro	gemini-2.5-pro	reasoning_tokens=27000	/
Gemini 2.5 Flash	gemini-2.5-flash	Default	/
DeepSeek R1	deepseek-reasoner	Default	/
Claude Opus 4	claude-opus-4-20250514	reasoning_tokens=27000	/
GPT-5 (minimal)	gpt5-2025-08-07	minimal*	/
Llama-4 Maverick	Llama-4-Maverick-17B-128E-Instruct-FP8	/	/
GPT-4o	chatgpt-4o-latest	/	/

* GPT-5 (minimal) sets `reasoning_effort=minimal`, which means zero reasoning tokens used.

Table 5: API configurations used in our evaluation.

A.4 API usage statistics

All evaluations are conducted using the standard API access provided by each company, ensuring scalability and consistency across models. As shown in Table 6, for each challenge, we perform five independent runs, and then calculate the average statistics per run, including number of input tokens, cached input tokens, reasoning tokens, response tokens, which adds up to total tokens.

The average cost per run (in USD) is determined by recorded tokens in each category and their respective pricing (in the unit of USD per million tokens). The input cost is computed by charging cached input tokens at the cached input rate and the remaining input tokens at the standard input rate. The output cost is computed from response tokens and reasoning tokens at the output rate. These two components then sum up to the total cost.

When models are augmented with tools, both input tokens and cached input tokens increase significantly. This is due to tool use inducing a multi-turn interaction: each round produces output that is fed back as input to the next round. Repeated context like prior turns is served from cache, inflating the cached input token counts. By contrast, reasoning tokens and output tokens in the table reflect only the final round, so intermediate generation across earlier turns appears on the input side rather than as additional output.

Model Name	Input		Cached Input		Reasoning		Response		Total	
	Tokens	Price (\$/1M)	Tokens	Price (\$/1M)	Tokens	Price (\$/1M)	Tokens	Price (\$/1M)	Tokens	Cost/run (\$)
GPT-5 (high, code & web)	330569	1.25	291050	0.125	10600	10.0	2286	10.0	343455	0.578
GPT-5 (high, code)	17728	1.25	16539	0.125	11379	10.0	1996	10.0	31102	0.158
GPT-5 (high)	864	1.25	55	0.125	24686	10.0	2035	10.0	27585	0.268
o3 (high)	864	2.0	200	0.5	17356	8.0	1322	8.0	19653	0.151
o4-mini (high)	864	1.1	140	0.275	16096	4.4	1536	4.4	18496	0.079
Gemini 2.5 Pro	1012	1.25	0	0.31	20943	10.0	3196	10.0	25151	0.243
Gemini 2.5 Flash	1012	0.3	0	0.075	22031	2.5	3726	2.5	26769	0.065
DeepSeek R1	821	0.56	320	0.07	19656	1.68	1236	1.68	21713	0.036
Claude Opus 4	1005	15.0	0	1.5	4566	75.0	8637	75.0	14208	1.005
GPT-5 (minimal)	864	1.25	24	0.125	0	10.0	2315	10.0	3179	0.024
GPT-4o	865	2.5	216	1.25	0	10.0	950	10.0	1815	0.012
Llama-4 Maverick	818	0.27	0	0.0	0	0.85	1383	0.85	2201	0.001

Table 6: Average token usage, pricing and cost across models in the evaluation of *CritPt* challenges.

A.5 Detailed analysis on example challenge: Quantum error detection

A.5.1 Design idea from expert

Quantum error correction (QEC) — the science of how to scalably suppress quantum noise in digital quantum computers — is a rapidly developing field that is becoming increasingly important to the development of quantum computing hardware. A key goal of QEC is to find QEC codes, protocols for encoding many physical qubits into logical qubits, with desirable properties, such as low time and space overhead and large error suppression. In this problem, we analyze some properties of the simplest possible QEC code, a small error detection code that can only detect errors and not correct them. Given this code’s small size, its properties can be worked out analytically, making it a useful test bed for understanding conceptually properties of more complicated large-scale QEC codes that can only be studied with numerical simulations, such as Monte Carlo sampling. This problem illustrates a set of questions a QEC researcher interested in this QEC code might ask to better understand its practical performance for a specific task, preparing a logical quantum state encoded in the code.

In subproblem 1, we describe a logical state preparation protocol for a specific quantum state in the QEC code and ask what the *physical* fidelity for that protocol is in terms of the error rate p in an idealized quantum computer with noisy gates. This quantity measures how much the quantum state is affected by the noise on the quantum computer without the help of error detection. It is a simple calculation that can be done by hand that requires knowing a little bit about the properties of stabilizers and Pauli matrices, common and well-known topics in quantum information and QEC.

In subproblem 2, we perform a similar calculation but now instead compute the *logical* fidelity of the logically encoded quantum state. This quantity measures how much the logical quantum state is affected by noise, and should be lower than the physical fidelity if the QEC code is working well. Comparing the logical and physical fidelities of operations in a QEC code is key to understanding the performance of the QEC code, so is a calculation of interest to researchers in the field. This calculation is a bit more involved than the previous subproblem, but still only relies on understanding the mathematics of stabilizers and Pauli matrices though in a more complicated setting. With patience this problem can be worked out by hand, but it is easier to solve it by writing some simple code that performs some combinatorics. The combinatorial calculation is necessary to understand how all of the different possible quantum gate errors lead to logical errors in the QEC code.

In subproblem 3, we now consider a different logical state preparation protocol for the same QEC code and ask for the logical state fidelity for that protocol. This state preparation protocol is significantly more complicated than the previous one. To solve this problem, one needs to write code to perform the combinatorics. Moreover, this code needs to be much more complicated than that required for subproblem 1. In addition to performing counting, it also needs to simulate how gate errors propagate through a quantum circuit, which involves running a non-trivial algorithm. Being able to answer problems such as this quickly, which can be stated quite simply but whose solution involves complex multi-step reasoning and code development, would be quite valuable to QEC researchers looking to quickly test many QEC codes.

A.5.2 Expert feedback on model responses

- **Subproblem 1:** Many models are able to solve this subproblem. Essentially all models understand the problem setup and the logic of how to obtain a solution. There is a sharp divide between models that are able to solve the problem correctly (e.g., o3 (high), o4-mini (high), GPT-5 (high), DeepSeek R1, Gemini 2.5 Pro) and those that are not. The incorrect models tend to immediately fail in the first steps of the calculation or make an illogical claim without support that leads to the wrong answer. Sometimes, the output of the model is long, convoluted, and horribly formatted, making it nearly impossible to tell where a logical mistake could have been made. Even among the correct models, there is quite a bit of separation in the quality of answers. o3 (high) produces some of the clearest, simple, and easy-to-understand derivations of the final solution. Other models, such as GPT-5 (high), o4-mini (high), and Gemini 2.5 Pro, produce answers with poor formatting (such as excessive use of bullet points) that are unnecessarily difficult to understand. These answers would be less useful for researchers, since it would make it difficult for them to verify the correctness of the result. DeepSeek R1 and Gemini 2.5 Pro appear to be at the edge of being able to solve the problem. They often answer correctly, but occasionally make an algebra mistake during the calculation. Other models, such as Gemini 2.5 Flash, GPT-4o, Llama-4 Maverick, would get lost in the derivation and often resort to

guessing answers. The GPT-5 (high, code) and GPT-5 (high, code & search) write code to solve the problem, which works but is completely unnecessary as this problem can be solved with simple counting that can be done by hand.

- **Subproblem 2:** In this subproblem, almost all of the models understand the high-level idea, which involves counting up different possible errors, assigning them probabilities, and combining the probabilities together in the correct way. However, most models are unable to execute the details of the calculation correctly. For the models able to solve the problem, when they give the correct answer, their explanations are generally clear, concise, and readable, making them useful for a researcher to understand the problem. When these models give incorrect answers, the derivations tend to become more vague, unclear, and poorly formatted. For the best performing models, when they give an incorrect answer the output tends to be well formatted and plausible, with more subtle errors hidden in the algebra or assumptions. This makes these models a bit dangerous to use, since their output can appear reasonable on their face but actually be completely wrong. The models that always give incorrect answers tend to have either minimal or confusing logic and produce wildly different solutions in each attempt, suggesting that they are guessing. At least in these models, it is quite apparent that the solutions are unreliable, which makes them less dangerous than the sometimes subtly-incorrect models. The GPT-5 (high, code) can reliably solve this problem. They write and execute code, which is not strictly necessary to solve the problem, but is a reasonable approach that works. However, despite giving consistently correct answers, all GPT-5's solutions are formatted quite poorly, often as a list of bullet points with incomplete sentences or vague comments. In principle, a researcher could decipher the material, but it certainly would be more useful if it was presented in a cleaner format, such as one might find in course lecture notes or a textbook.
- **Subproblem 3:** All models except for GPT-5 (high, code) and GPT-5 (high, code & search) completely fail on this subproblem. This subproblem requires using the same general reasoning as subproblem 2, but now in a more complicated setting that involves writing and executing code to perform the counting of quantum errors. All other models (including GPT-5 (high)) perform rampant guessing, inventing the numbers needed to get a final solution. The best incorrect models tend to produce answers that are some combination of vague, confusing, poorly formatted, or excessively short. The worst incorrect models simply repeat the information presented in the prompt and guess an answer (such as 1). Most incorrect models produce different random answers in each trial. In all cases, it would be easy for a researcher to see that these models are unable to give a correct solution. However, in no case did the incorrect models indicate that they did not know how to solve the problem, instead giving long and complicated justifications for clearly incorrect solutions. Even though GPT-5 (high, code) could fairly consistently produce the correct answer, its solutions were again poorly formatted (e.g., everything was a list of vague bullet points) making them much less useful than they could be. The poor formatting tends to make me trust the model less, since I have seen poor formatting coincide with incorrect logic in other models. Therefore, even though GPT-5 can often get the right answer here, its presentation to the user can still be significantly improved to make it more trustworthy for research use cases.

References

- [1] P. W. Anderson. More is different: broken symmetry and the nature of the hierarchical structure of science. *Science*, 177(4047):393–396, 1972.
- [2] R. Sinatra, P. Deville, M. Szell, D. Wang, and A.-L. Barabási. A century of physics. *Nature Physics*, 11(10):791–796, 2015.
- [3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. pages, 4171–4186, 2019.
- [5] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [6] F. M. Delgado-Chaves, M. J. Jennings, A. Atalaia, J. Wolff, R. Horvath, Z. M. Mamdouh, J. Baumbach, and L. Baumbach. Transforming literature screening: The emerging role of large language models in systematic reviews. *Proceedings of the National Academy of Sciences*, 122(2): e2411962122, 2025.
- [7] D. Scherbakov, N. Hubig, V. Jansari, A. Bakumenko, and L. A. Lenert. The emergence of large language models as tools in literature reviews: a large language model-assisted systematic review. *Journal of the American Medical Informatics Association*, 32(6):1071–1086, 2025.
- [8] S. Pramanick, R. Chellappa, and S. Venugopalan. SPIQA: A dataset for multimodal question answering on scientific papers. *Advances in Neural Information Processing Systems*, 37:118807–118833, 2024.
- [9] T. Gao, H. Yen, J. Yu, and D. Chen. Enabling large language models to generate text with citations. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [10] Y. Wang, Q. Guo, W. Yao, H. Zhang, X. Zhang, Z. Wu, M. Zhang, X. Dai, Q. Wen, W. Ye, et al. Autosurvey: Large Language Models can automatically write surveys. *Advances in neural information processing systems*, 37:115119–115145, 2024.
- [11] A. Asai, J. He, R. Shao, W. Shi, A. Singh, J. C. Chang, K. Lo, L. Soldaini, S. Feldman, M. D’arcy, et al. Openscholar: Synthesizing scientific literature with retrieval-augmented LMs. *arXiv preprint arXiv:2411.14199*, 2024.
- [12] M. D. Skarlinski, S. Cox, J. M. Laurent, J. D. Braza, M. Hinks, M. J. Hammerling, M. Ponnappati, S. G. Rodrigues, and A. D. White. Language agents achieve superhuman synthesis of scientific knowledge. *arXiv preprint arXiv:2409.13740*, 2024.
- [13] H. Cui, Z. Shamsi, G. Cheon, X. Ma, S. Li, M. Tikhonovskaya, P. C. Norgaard, N. Mudur, M. B. Plomecka, P. Raccuglia, et al. CURIE: evaluating LLMs on multitask scientific long-context understanding and reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [14] OpenAI. Introducing GPT-5, 2025. <https://openai.com/index/introducing-gpt-5/>.
- [15] OpenAI. Introducing OpenAI o3 and o4-mini, 2025. <https://openai.com/index/introducing-o3-and-o4-mini/>.
- [16] Google. Gemini 2.5: Our most intelligent AI model, 2025. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/#gemini-2-5-thinking>.
- [17] D. Guo, D. Yang, H. Zhang, J. Song, P. Wang, Q. Zhu, R. Xu, R. Zhang, S. Ma, X. Bi, et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081): 633–638, 2025.
- [18] Anthropic. Introducing Claude 4, 2025. <https://www.anthropic.com/news/claude-4>.

- [19] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford, et al. GPT-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- [20] Meta. The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation, 2025. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>.
- [21] A. Jaech, A. Kalai, A. Lerer, A. Richardson, A. El-Kishky, A. Low, A. Helyar, A. Madry, A. Beutel, A. Carney, et al. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- [22] G. Comanici, E. Bieber, M. Schaekermann, I. Pasupat, N. Sachdeva, I. Dhillon, M. Blistein, O. Ram, D. Zhang, E. Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- [23] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- [24] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- [25] T. Schick, J. Dwivedi-Yu, R. Dessí, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom. Toolformer: language models can teach themselves to use tools. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023.
- [26] X. Wang, Y. Chen, L. Yuan, Y. Zhang, Y. Li, H. Peng, and H. Ji. Executable code actions elicit better LLM agents. In *Proceedings of the International Conference on Machine Learning*, 2024.
- [27] L. Yuan, Y. Chen, X. Wang, Y. Fung, H. Peng, and H. Ji. CRAFT: Customizing LLMs by creating and retrieving from specialized toolsets. In *The Twelfth International Conference on Learning Representations*, 2024.
- [28] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [29] X. Wang, J. Wei, D. Schuurmans, Q. V. Le, E. H. Chi, S. Narang, A. Chowdhery, and D. Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023.
- [30] C. V. Snell, J. Lee, K. Xu, and A. Kumar. Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [31] R. Hazra, G. Venturato, P. Z. Dos Martires, and L. De Raedt. Have large language models learned to reason? a characterization via 3-sat. In *Second Conference on Language Modeling*, 2025.
- [32] K. Gandhi, A. K. Chakravarthy, A. Singh, N. Lile, and N. Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective STaRs. In *Second Conference on Language Modeling*, 2025.
- [33] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [34] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [35] C. E. Jimenez, J. Yang, A. Wettig, S. Yao, K. Pei, O. Press, and K. R. Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *The Twelfth International Conference on Learning Representations*, 2024.

- [36] Google DeepMind. Advanced version of Gemini with deep think officially achieves gold-medal standard at the International Mathematical Olympiad, 2025.
- [37] A. El-Kishky, A. Wei, A. Saraiva, B. Minaiev, D. Selsam, D. Dohan, F. Song, H. Lightman, I. Clavera, J. Pachocki, et al. Competitive programming with large reasoning models. *arXiv preprint arXiv:2502.06807*, 2025.
- [38] M. Balunović, J. Dekoninck, I. Petrov, N. Jovanović, and M. Vechev. MathArena: Evaluating llms on uncontaminated math competitions. *arXiv preprint arXiv:2505.23281*, 2025.
- [39] C. He, R. Luo, Y. Bai, S. Hu, Z. Thai, J. Shen, J. Hu, X. Han, Y. Huang, Y. Zhang, et al. OlympiadBench: A challenging benchmark for promoting AGI with Olympiad-level bilingual multimodal scientific problems. pages, 3828–3850, 2024.
- [40] N. Jain, K. Han, A. Gu, W.-D. Li, F. Yan, T. Zhang, S. Wang, A. Solar-Lezama, K. Sen, and I. Stoica. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [41] S. Qiu, S. Guo, Z.-Y. Song, Y. Sun, Z. Cai, J. Wei, T. Luo, Y. Yin, H. Zhang, Y. Hu, et al. PHYBench: Holistic evaluation of physical perception and reasoning in large language models. *arXiv preprint arXiv:2504.16074*, 2025.
- [42] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [43] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021.
- [44] Y. Wang, X. Ma, G. Zhang, Y. Ni, A. Chandra, S. Guo, W. Ren, A. Arulraj, X. He, Z. Jiang, et al. MMLU-Pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290, 2024.
- [45] X. Wang, Z. Hu, P. Lu, Y. Zhu, J. Zhang, S. Subramaniam, A. R. Loomba, S. Zhang, Y. Sun, and W. Wang. SciBench: Evaluating college-level scientific problem-solving abilities of large language models. pages, 50622–50649. PMLR, 2024.
- [46] X. Xu, Q. Xu, T. Xiao, T. Chen, Y. Yan, J. ZHANG, S. Diao, C. Yang, and Y. Wang. UGPhysics: A comprehensive benchmark for undergraduate physics reasoning with large language models. In *Forty-second International Conference on Machine Learning*, 2025.
- [47] D. Rein, B. L. Hou, A. C. Stickland, J. Petty, R. Y. Pang, J. Dirani, J. Michael, and S. R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
- [48] M. Tian, L. Gao, S. Zhang, X. Chen, C. Fan, X. Guo, R. Haas, P. Ji, K. Krongchon, Y. Li, et al. Scicode: A research coding benchmark curated by scientists. *Advances in Neural Information Processing Systems*, 37:30624–30650, 2024.
- [49] E. Glazer, E. Erdil, T. Besiroglu, D. Chicharro, E. Chen, A. Gunning, C. F. Olsson, J.-S. Denain, A. Ho, E. d. O. Santos, et al. FrontierMath: a benchmark for evaluating advanced mathematical reasoning in AI. *arXiv preprint arXiv:2411.04872*, 2024.
- [50] D. J. Chung, Z. Gao, Y. Kvsiuk, T. Li, M. Münchmeyer, M. Rudolph, F. Sala, and S. C. Tadepalli. Theoretical physics benchmark (TPBench)—a dataset and study of AI reasoning capabilities in theoretical physics. *arXiv preprint arXiv:2502.15815*, 2025.
- [51] L. Phan, A. Gatti, Z. Han, N. Li, J. Hu, H. Zhang, C. B. C. Zhang, M. Shaaban, J. Ling, S. Shi, et al. Humanity’s last exam. *arXiv preprint arXiv:2501.14249*, 2025.
- [52] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, P. Chandak, S. Liu, P. Van Katwyk, A. Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.

- [53] Z. Wu, L. Qiu, A. Ross, E. Akyürek, B. Chen, B. Wang, N. Kim, J. Andreas, and Y. Kim. Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. pages, 1819–1862, 2024.
- [54] N. Balepur, A. Ravichander, and R. Rudinger. Artifacts or abduction: How do LLMs answer multiple-choice questions without the question? pages, 10308–10330, 2024.
- [55] C. Deng, Y. Zhao, X. Tang, M. Gerstein, and A. Cohan. Investigating data contamination in modern benchmarks for large language models. pages, 8698–8711, 2024.
- [56] S. Ott, A. Barbosa-Silva, K. Blagec, J. Brauner, and M. Samwald. Mapping global dynamics of benchmark creation and saturation in artificial intelligence. *Nature Communications*, 13(1):6793, 2022.
- [57] Y. Li, Y. Guo, F. Guerin, and C. Lin. An open-source data contamination report for large language models. pages, 528–541, 2024.
- [58] N. Balepur, R. Rudinger, and J. L. Boyd-Graber. Which of these best describes multiple choice evaluation with LLMs? A) forced B) flawed C) fixable D) all of the above. pages, 3394–3418. Association for Computational Linguistics, 2025.
- [59] J. Dodge, M. Sap, A. Marasović, W. Agnew, G. Ilharco, D. Groeneveld, M. Mitchell, and M. Gardner. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. pages, 1286–1305, 2021.
- [60] S. Golchin and M. Surdeanu. Time travel in LLMs: Tracing data contamination in large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [61] M. Roberts, H. Thakur, C. Herlihy, C. White, and S. Dooley. To the cutoff... and beyond? a longitudinal perspective on llm data contamination. In *The Twelfth International Conference on Learning Representations*, 2023.
- [62] P. Wang, L. Li, L. Chen, Z. Cai, D. Zhu, B. Lin, Y. Cao, L. Kong, Q. Liu, T. Liu, et al. Large language models are not fair evaluators. pages, 9440–9450, 2024.
- [63] J. Ye, Y. Wang, Y. Huang, D. Chen, Q. Zhang, N. Moniz, T. Gao, W. Geyer, C. Huang, P.-Y. Chen, et al. Justice or prejudice? quantifying biases in LLM-as-a-judge. In *International Conference on Learning Representations*, 2025.
- [64] M. T. R. Laskar, S. Alqahtani, M. S. Bari, M. Rahman, M. A. M. Khan, H. Khan, I. Jahan, A. Bhuiyan, C. W. Tan, M. R. Parvez, E. Hoque, S. Joty, and J. Huang. A systematic survey and critical review on evaluating large language models: Challenges, limitations, and recommendations. pages, 13785–13816. Association for Computational Linguistics, 2024.
- [65] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, Kumar, et al. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3:e103, 2017.
- [66] PhySH – Physics Subject Headings. <https://physh.org/about>. Accessed: August 18, 2025.
- [67] A. T. Kalai, O. Nachum, S. S. Vempala, and E. Zhang. Why language models hallucinate. *arXiv preprint arXiv:2509.04664*, 2025.
- [68] D. Gottesman. Quantum fault tolerance in small experiments. *arXiv preprint arXiv:1610.03507*, 2016.
- [69] C. Vuillot. Is error detection helpful on IBM 5Q chips? *Quantum Inf. Comput.*, 18(11):0949, 2017.
- [70] N. M. Linke, M. Gutierrez, K. A. Landsman, C. Figgatt, S. Debnath, K. R. Brown, and C. Monroe. Fault-tolerant quantum error detection. *Sci. Adv.*, 3(10):e1701074, 2017.
- [71] R. Harper and S. T. Flammia. Fault-tolerant logical gates in the IBM quantum experience. *Phys. Rev. Lett.*, 122:080504, 2019.

- [72] P. Komar, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin. A quantum network of clocks. *Nature Physics*, 10(8):582–587, 2014.
- [73] Z. Zhang and Q. Zhuang. Distributed quantum sensing. *Quantum Science and Technology*, 6(4):043001, 2021.
- [74] A. Zang, A. Kolar, A. Gonzales, J. Chung, S. K. Gray, R. Kettimuthu, T. Zhong, and Z. H. Saleem. Quantum advantage in distributed sensing with noisy quantum networks. *arXiv preprint arXiv:2409.17089*, 2024.
- [75] A. Zang, T.-X. Zheng, P. C. Maurer, F. T. Chong, M. Suchara, and T. Zhong. Enhancing noisy quantum sensing by GHZ state partitioning. *arXiv preprint arXiv:2507.02829*, 2025.
- [76] A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, 1981.
- [77] A. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393, 1982.
- [78] A. Albrecht and P. J. Steinhardt. Cosmology for grand unified theories with radiatively induced symmetry breaking. *Phys. Rev. Lett.*, 48:1220–1223, 1982.
- [79] A. Linde. Chaotic inflation. *Physics Letters B*, 129(3):177–181, 1983.
- [80] K. Freese, J. A. Frieman, and A. V. Olinto. Natural inflation with pseudo Nambu-Goldstone bosons. *Phys. Rev. Lett.*, 65:3233–3236, 1990.
- [81] W. H. Kinney and K. T. Mahanthappa. Natural inflation from fermion loops. *Phys. Rev. D*, 52:5529–5537, 1995.
- [82] N. Arkani-Hamed, H.-C. Cheng, P. Creminelli, and L. Randall. Pseudonatural inflation. *Journal of Cosmology and Astroparticle Physics*, 2003(07):003, 2003.
- [83] P. Adshead and M. Wyman. Natural inflation on a steep potential with classical non-abelian gauge fields. *Phys. Rev. Lett.*, 108:261302, 2012.
- [84] P. Adshead and M. Wyman. Gauge-flation trajectories in chromo-natural inflation. *Phys. Rev. D*, 86:043530, 2012.
- [85] C. Long, L. McAllister, and P. McGuirk. Aligned natural inflation in string theory. *Phys. Rev. D*, 90:023501, 2014.
- [86] A. Maleknejad. Gravitational leptogenesis in axion inflation with SU(2) gauge field. *Journal of Cosmology and Astroparticle Physics*, 2016(12):027, 2016.
- [87] T. Fujita, K. Murai, I. Obata, and M. Shiraishi. Gravitational wave trispectrum in the axion-SU(2) model. *Journal of Cosmology and Astroparticle Physics*, 2022(01):007, 2022.
- [88] V. Gluscevic and M. Kamionkowski. Testing parity-violating mechanisms with cosmic microwave background experiments. *Phys. Rev. D*, 81:123529, 2010.
- [89] W. J. Wolf. Minimizing the tensor-to-scalar ratio in single-field inflation models. *Phys. Rev. D*, 110:043521, 2024.
- [90] S. Kachru, R. Kallosh, A. Linde, J. Maldacena, L. McAllister, and S. P. Trivedi. Towards inflation in string theory. *Journal of Cosmology and Astroparticle Physics*, 2003(10):013, 2003.
- [91] E. Witten. Some properties of the $(\psi\psi)_2$ model in two dimensions. *Nuclear Physics B*, 142(3):285–300, 1978.
- [92] Y. Y. Goldschmidt. A Kosterlitz-Thouless phase transition associated with the supersymmetric sine-gordon theory. *Nuclear Physics B*, 270:29–38, 1986.
- [93] G. Moore and N. Read. Nonabelions in the fractional quantum Hall effect. *Nuclear Physics B*, 360(2-3):362–396, 1991.

- [94] M. Milovanović and N. Read. Edge excitations of paired fractional quantum Hall states. *Physical Review B*, 53(20):13559, 1996.
- [95] N. Schiller, B. A. Katzir, A. Stern, E. Berg, N. H. Lindner, and Y. Oreg. Superconductivity and fermionic dissipation in quantum Hall edges. *Physical Review B*, 107(16):L161105, 2023.
- [96] J. Cao, A. Kou, and E. Fradkin. Signatures of parafermion zero modes in fractional quantum Hall–superconductor heterostructures. *Physical Review B*, 109(16):L161106, 2024.
- [97] J. May-Mann, A. Stern, and T. Devakul. Theory of half-integer fractional quantum spin hall insulator edges. *arXiv preprint arXiv:2403.03964*, 2024.
- [98] A. Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006.
- [99] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma. Non-abelian anyons and topological quantum computation. *Reviews of Modern Physics*, 80(3):1083–1159, 2008.
- [100] M. B. Hastings, C. Nayak, and Z. Wang. Metaplectic anyons, majorana zero modes, and their computational power. *Physical Review B—Condensed Matter and Materials Physics*, 87(16):165421, 2013.
- [101] D. J. Clarke, J. Alicea, and K. Shtengel. Exotic non-abelian anyons from conventional fractional quantum Hall states. *Nature communications*, 4(1):1348, 2013.
- [102] N. H. Lindner, E. Berg, G. Refael, and A. Stern. Fractionalizing Majorana fermions: Non-abelian statistics on the edges of abelian quantum Hall states. *Physical Review X*, 2(4):041002, 2012.
- [103] M. Cheng. Superconducting proximity effect on the edge of fractional topological insulators. *Physical Review B—Condensed Matter and Materials Physics*, 86(19):195126, 2012.
- [104] M. Barkeshli, C.-M. Jian, and X.-L. Qi. Theory of defects in abelian topological states. *Physical Review B—Condensed Matter and Materials Physics*, 88(23):235103, 2013.
- [105] P. Fendley. Parafermionic edge zero modes in Zn-invariant spin chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(11):P11020, 2012.
- [106] E. Verlinde. Fusion rules and modular transformations in 2d conformal field theory. *Nuclear Physics B*, 300:360–376, 1988.
- [107] J. Fröhlich, J. Fuchs, I. Runkel, and C. Schweigert. Kramers-Wannier duality from conformal defects. *Physical review letters*, 93(7):070601, 2004.
- [108] J. Fröhlich, J. Fuchs, I. Runkel, and C. Schweigert. Duality and defects in rational conformal field theory. *Nuclear Physics B*, 763(3):354–430, 2007.
- [109] P. Francesco, P. Mathieu, and D. Sénéchal. *Conformal field theory*. Springer Science & Business Media, 2012.
- [110] C.-M. Chang, Y.-H. Lin, S.-H. Shao, Y. Wang, and X. Yin. Topological defect lines and renormalization group flows in two dimensions. *Journal of High Energy Physics*, 2019(1):1–85, 2019.
- [111] Z.-M. Huang, L. Colmenarez, M. Müller, and S. Diehl. Coherent information as a mixed-state topological order parameter of fermions. *arXiv preprint arXiv:2412.12279*, 2024.
- [112] R. Fan, Y. Bao, E. Altman, and A. Vishwanath. Diagnostics of mixed-state topological order and breakdown of quantum memory. *PRX Quantum*, 5(2):020343, 2024.
- [113] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002.
- [114] H. Nishimori. Number. 111. Clarendon Press, 2001.
- [115] M. P. Zaletel, R. S. Mong, and F. Pollmann. Flux insertion, entanglement, and quantized responses. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(10):P10007, 2014.

- [116] J. I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete. Matrix product states and projected entangled pair states: Concepts, symmetries, theorems. *Reviews of Modern Physics*, 93(4):045003, 2021.
- [117] Z.-M. Huang, S. Diehl, and X.-Q. Sun. Topological response in open quantum systems with weak symmetries. *arXiv preprint arXiv:2504.02941*, 2025.
- [118] K. Sun, Z. Gu, H. Katsura, and S. Das Sarma. Nearly flatbands with nontrivial topology. *Physical review letters*, 106(23):236803, 2011.
- [119] T. Neupert, L. Santos, C. Chamon, and C. Mudry. Fractional quantum Hall states at zero magnetic field. *Physical review letters*, 106(23):236804, 2011.
- [120] P. Mai, J. Zhao, B. E. Feldman, and P. W. Phillips. $1/4$ is the new $1/2$ when topology is intertwined with Motttness. *Nature communications*, 14(1):5999, 2023.
- [121] P. Mai, B. E. Feldman, and P. W. Phillips. Topological Mott insulator at quarter filling in the interacting haldane model. *Physical Review Research*, 5(1):013162, 2023.
- [122] P. Mai, J. Zhao, T. A. Maier, B. Bradlyn, and P. W. Phillips. Topological phase transition without single particle gap closing in strongly correlated systems. *Physical Review B*, 110(7):075105, 2024.
- [123] W. Fu and S. Sachdev. Numerical study of fermion and boson models with infinite-range random interactions. *Phys. Rev. B*, 94:035135, 2016.
- [124] J. Maldacena and D. Stanford. Remarks on the Sachdev-Ye-Kitaev model. *Phys. Rev. D*, 94:106002, 2016.
- [125] T. Izubuchi, X. Ji, L. Jin, I. W. Stewart, and Y. Zhao. Factorization theorem relating Euclidean and light-cone parton distributions. *Phys. Rev. D*, 98(5):056004, 2018.
- [126] S. Moch, J. A. M. Vermaseren, and A. Vogt. The three loop splitting functions in QCD: The nonsinglet case. *Nucl. Phys. B*, 688:101–134, 2004.
- [127] Y. Su, J. Holligan, X. Ji, F. Yao, J.-H. Zhang, and R. Zhang. Resumming quark’s longitudinal momentum logarithms in LaMET expansion of lattice PDFs. *Nucl. Phys. B*, 991:116201, 2023.
- [128] X. Ji, Y. Liu, Y.-S. Liu, J.-H. Zhang, and Y. Zhao. Large-momentum effective theory. *Reviews of Modern Physics*, 93(3):035005, 2021.
- [129] X. Ji. Parton physics on a euclidean lattice. *Physical Review Letters*, 110(26):262002, 2013.
- [130] X. Ji. Parton physics from large-momentum effective field theory. *Sci. China Phys. Mech. Astron.*, 57:1407–1412, 2014.
- [131] K. Horodecki, M. Horodecki, P. Horodecki, and J. Oppenheim. General paradigm for distilling classical key from quantum states. *IEEE Transactions on Information Theory*, 55(4):1898–1929, 2009.
- [132] G. Smith and P. Wu. Additivity of quantum capacities in simple non-degradable quantum channels. *IEEE Transactions on Information Theory*, 2025.
- [133] A. Lesniewski and M. B. Ruskai. Monotone Riemannian metrics and relative entropy on noncommutative probability spaces. *Journal of Mathematical Physics*, 40(11):5702–5724, 1999.
- [134] F. Hiai and M. B. Ruskai. Contraction coefficients for noisy quantum channels. *Journal of Mathematical Physics*, 57(1), 2016.
- [135] T. M. Hoang, Y. Ma, J. Ahn, J. Bang, F. Robicheaux, Z.-Q. Yin, and T. Li. Torsional optomechanics of a levitated nonspherical nanoparticle. *Physical review letters*, 117(12):123604, 2016.
- [136] G. Zhang, H. Zhang, and Z.-q. Yin. Scalable universal quantum gates between nitrogen-vacancy centers in levitated nanodiamonds arrays. *arXiv preprint arXiv:2504.08194*, 2025.

- [137] J. Rieser, M. A. Ciampini, H. Rudolph, N. Kiesel, K. Hornberger, B. A. Stickler, M. Aspelmeyer, and U. Delić. Tunable light-induced dipole-dipole interaction between optically levitated nanoparticles. *Science*, 377(6609):987–990, 2022.
- [138] M. Manceau, F. Khalili, and M. Chekhova. Improving the phase super-sensitivity of squeezing-assisted interferometers by squeeze factor unbalancing. *New Journal of Physics*, 19(1):013014, 2017.
- [139] R. Nehra, R. Sekine, L. Ledezma, Q. Guo, R. M. Gray, A. Roy, and A. Marandi. Few-cycle vacuum squeezing in nanophotonics. *Science*, 377(6612):1333–1337, 2022.
- [140] K. Murase, F. Oikonomou, and M. Petropoulou. Blazar flares as an origin of high-energy cosmic neutrinos? *Astrophys. J.*, 865(2):124, 2018.
- [141] P. Padovani, F. Oikonomou, M. Petropoulou, P. Giommi, and E. Resconi. TXS 0506+056, the first cosmic neutrino source, is not a BL Lac. *Mon. Not. Roy. Astron. Soc.*, 484(1):L104–L108, 2019.
- [142] T. H. Boyer. Quantum electromagnetic zero-point energy of a conducting spherical shell and the casimir model for a charged particle. *Physical Review*, 174(5):1764, 1968.
- [143] L. S. Brown and G. Gabrielse. Geonium theory: Physics of a single electron or ion in a Penning trap. *Reviews of Modern Physics*, 58(1):233, 1986.
- [144] L. S. Brown, K. Helmersen, and J. Tan. Cyclotron motion in a spherical microwave cavity. *Physical Review A*, 34(4):2638, 1986.
- [145] G. Barton and N. S. Fawcett. Quantum electromagnetics of an electron near mirrors. *Physics Reports*, 170(1):1–95, 1988.
- [146] X. Fan, T. Myers, B. Sukra, and G. Gabrielse. Measurement of the electron magnetic moment. *Physical review letters*, 130(7):071801, 2023.
- [147] M. Kitagawa and M. Ueda. Squeezed spin states. *Physical Review A*, 47(6):5138, 1993.
- [148] D. J. Wineland, J. J. Bollinger, W. M. Itano, and D. J. Heinzen. Squeezed atomic states and projection noise in spectroscopy. *Physical Review A*, 50(1):67, 1994.
- [149] J. Ma, X. Wang, and F. Nori. Quantum spin squeezing. *Physics Reports*, 509(2-3):89–165, 2011.
- [150] A. Chu, P. He, J. K. Thompson, and A. M. Rey. Quantum enhanced cavity QED interferometer with partially delocalized atoms in lattices. *Physical Review Letters*, 127(21):210401, 2021.
- [151] D. Barberena, A. Chu, J. K. Thompson, and A. M. Rey. Trade-offs between unitary and measurement induced spin squeezing in cavity QED. *Physical Review Research*, 6(3):L032037, 2024.
- [152] D. Goluskin. *Internally heated convection and Rayleigh-Bénard convection*. Springer, 2016.
- [153] C. Liu, M. Sharma, K. Julien, and E. Knobloch. Fixed-flux Rayleigh-Bénard convection in doubly periodic domains: generation of large-scale shear. *Journal of Fluid Mechanics*, 979:A19, 2024.
- [154] S. Chandrasekhar. *Hydrodynamic and hydromagnetic stability*. Courier Corporation, 2013.
- [155] F. H. Busse. On the stability of two-dimensional convection in a layer heated from below. *Journal of Mathematics and Physics*, 46(1-4):140–150, 1967.
- [156] P. Manneville. Rayleigh-Bénard convection: Thirty years of experimental, theoretical, and modeling work. pages, 41–65. Springer New York, New York, NY, 2006.
- [157] C. Liu and E. Knobloch. Single-mode solutions for convection and double-diffusive convection in porous media. *Fluids*, 7(12):373, 2022.
- [158] F. H. Busse. Non-linear properties of thermal convection. *Reports on Progress in Physics*, 41(12):1929, 1978.
- [159] D. A. Nield and A. Bejan. *Convection in porous media*. Springer, 2006.

- [160] D. Nield and C. T. Simmons. A brief introduction to convection in porous media. *Transport in Porous Media*, 130(1):237–250, 2019.
- [161] O. V. Trevisan and A. Bejan. Mass and heat transfer by high rayleigh number convection in a porous medium heated from below. *International Journal of Heat and Mass Transfer*, 30(11): 2341–2356, 1987.
- [162] D. Hewitt. Vigorous convection in porous media. *Proceedings of the Royal Society A*, 476(2239): 20200111, 2020.
- [163] J. F. Beacom, N. F. Bell, and G. D. Mack. General upper bound on the dark matter total annihilation cross section. *Phys. Rev. Lett.*, 99:231301, 2007.
- [164] A. Alenezi, C. Cesarotti, S. Gori, and J. Shelton. Discovery prospects for a minimal dark matter model at cosmic and intensity frontier experiments. 2025.
- [165] J. P. Bartolotta, S. B. Jäger, J. T. Reilly, M. A. Norcia, J. K. Thompson, G. Smith, and M. J. Holland. Entropy transfer from a quantum particle to a classical coherent light field. *Physical Review Research*, 4(1):013218, 2022.
- [166] R. Bellman and R. S. Lehman. The reciprocity formula for multidimensional theta functions. *Proceedings of the American Mathematical Society*, 12(6):954–961, 1961.
- [167] A. Jeffrey and D. Zwillinger. *Table of integrals, series, and products*. Academic Press, 2007.
- [168] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić. Quantum scarred eigenstates in a Rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations. *Physical Review B*, 98(15), 2018.
- [169] T. Zhou, A. Y. Guo, S. Xu, X. Chen, and B. Swingle. Hydrodynamic theory of scrambling in chaotic long-range interacting systems. *Physical Review B*, 107(1):014201, 2023.
- [170] O. Hallatschek and D. S. Fisher. Acceleration of evolutionary spread by long-range dispersal. *Proceedings of the National Academy of Sciences*, 111(46):E4911–E4919, 2014.
- [171] T. Zhou and A. Nahum. Emergent statistical mechanics of entanglement in random unitary circuits. *Physical Review B*, 99(17):174205, 2019.
- [172] S.-H. Chen and M. Kotlarchyk. *Interactions of photons and neutrons with matter*. World Scientific, 2007.
- [173] M. Tsang, R. Nair, and X.-M. Lu. Quantum theory of superresolution for two incoherent optical point sources. *Physical Review X*, 6(3):031033, 2016.
- [174] D. A. Meyer and T. G. Wong. Connectivity is a poor indicator of fast quantum search. *Physical review letters*, 114(11):110503, 2015.
- [175] B. Collins and P. Śniady. Integration with respect to the Haar measure on unitary, orthogonal and symplectic group. *Communications in Mathematical Physics*, 264(3):773–795, 2006.
- [176] J. T. Reilly, S. B. Jäger, J. D. Wilson, J. Cooper, S. Eggert, and M. J. Holland. Speeding up squeezing with a periodically driven dicke model. *Physical Review Research*, 6(3):033090, 2024.
- [177] J. D. Wilson, J. T. Reilly, H. Zhang, C. Luo, A. Chu, J. K. Thompson, A. M. Rey, and M. J. Holland. Entangled matter waves for quantum enhanced sensing. *Physical Review A*, 110(4): L041301, 2024.
- [178] S. B. Jäger, T. Schmit, G. Morigi, M. J. Holland, and R. Betzholtz. Lindblad master equations for quantum systems coupled to dissipative bosonic modes. *Physical Review Letters*, 129(6):063601, 2022.
- [179] Y. Zhang, X. Chen, and E. Chitambar. Building multiple access channels with a single particle. *Quantum*, 6:653, 2022.

- [180] S. Horvat and B. Dakić. Quantum enhancement to information acquisition speed. *New Journal of Physics*, 23(3):033008, 2021.
- [181] A. Mu, Z. Sun, and A. J. Millis. Adequacy of the dynamical mean field theory for low density and dirac materials. *Phys. Rev. B*, 109:115154, 2024.
- [182] A. Mu, Z. Sun, and A. J. Millis. Optical conductivity of the two-dimensional hubbard model: Vertex corrections, emergent galilean invariance, and the accuracy of the single-site dynamical mean field approximation. *Phys. Rev. B*, 106:085142, 2022.
- [183] A. Rosch and P. C. Howell. Zero-temperature optical conductivity of ultraclean fermi liquids and superconductors. *Phys. Rev. B*, 72:104510, 2005.
- [184] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *nature*, 415(6867):39–44, 2002.
- [185] I. Bloch, J. Dalibard, and S. Nascimbene. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8(4):267–276, 2012.
- [186] C. Gross and I. Bloch. Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355):995–1001, 2017.
- [187] C. Gross and W. S. Bakr. Quantum gas microscopy for single atom and spin detection. *Nature Physics*, 17(12):1316–1323, 2021.
- [188] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. *Science*, 377(6608):885–889, 2022.
- [189] I. H. Kim. Holographic quantum simulation. *arXiv:1702.02093*, 2017.
- [190] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter. Holographic quantum algorithms for simulating correlated spin systems. *Phys. Rev. Research*, 3:033002, 2021.
- [191] F. Barratt, J. Dborin, M. Bal, V. Stojevic, F. Pollmann, and A. G. Green. Parallel quantum simulation of large systems on small NISQ computers. *npj Quantum Inf.*, 7(1), 2021.
- [192] E. Chertkov, J. Bohnet, D. Francois, J. Gaebler, D. Gresh, A. Hankin, K. Lee, D. Hayes, B. Neyenhuis, R. Stutz, A. C. Potter, and M. Foss-Feig. Holographic dynamics simulations with a trapped-ion quantum computer. *Nat. Phys.*, 18:1074, 2022.
- [193] D. Niu, R. Haghshenas, Y. Zhang, M. Foss-Feig, G. K.-L. Chan, and A. C. Potter. Holographic simulation of correlated electrons on a trapped ion quantum processor. *arXiv:2112.10810*, 2021.
- [194] Y. Zhang, S. Jahanbani, D. Niu, R. Haghshenas, and A. C. Potter. Qubit-efficient simulation of thermal states with quantum tensor networks. *arXiv:2205.06299*, 2022.
- [195] M. DeCross, E. Chertkov, M. Kohagen, and M. Foss-Feig. Qubit-reuse compilation with mid-circuit measurement and reset. *arXiv:2210.08039*, 2022.
- [196] E. Chertkov and B. K. Clark. Computational inverse method for constructing spaces of quantum models from wave functions. *Phys. Rev. X*, 8:031029, 2018.
- [197] E. Chertkov, B. Villalonga, and B. K. Clark. Engineering topological models with a general-purpose symmetry-to-Hamiltonian approach. *Phys. Rev. Res.*, 2:023348, 2020.
- [198] X.-L. Qi and D. Ranard. Determining a local Hamiltonian from a single eigenstate. *Quantum*, 3:159, 2019. ISSN 2521-327X.
- [199] E. O. Powell. Growth rate and generation time of bacteria, with special reference to continuous culture. *Journal of General Microbiology*, 15(3):492–511, 1956.
- [200] F. Jafarpour, C. S. Wright, H. Gudjonson, J. Riebling, E. Dawson, K. Lo, A. Fiebig, S. Crosson, A. R. Dinner, and S. Iyer-Biswas. Bridging the timescales of single-cell and population dynamics. *Physical Review X*, 8(2):021007, 2018.

- [201] F. Barber, J. Min, A. W. Murray, and A. Amir. Modeling the impact of single-cell stochasticity and size control on the population growth rate in asymmetrically dividing cells. *PLOS Computational Biology*, 17(6):e1009080, 2021.
- [202] Y. Hein and F. Jafarpour. Asymptotic decoupling of population growth rate and cell size distribution. *Physical Review Research*, 6(4):043006, 2024.
- [203] E. Levien, Y. Heřn, and F. Jafarpour. Size-structured populations with growth fluctuations: Feynman–Kac formula and decoupling. *arXiv preprint arXiv:2508.14680*, 2025.
- [204] C. N. Hinshelwood. On the chemical kinetics of autotrophic systems. pages, 745–755, 1952.
- [205] R. M. Gray. Toeplitz and circulant matrices: A review. *Foundations and Trends in Communications and Information Theory*, 2(3):155–239, 2006.
- [206] Y. Hein and F. Jafarpour. Competition between transient oscillations and early stochasticity in exponentially growing populations. *Physical Review Research*, 6(3):033320, 2024.
- [207] T. Dauxois, F. Di Patti, D. Fanelli, and A. J. McKane. Enhanced stochastic oscillations in autocatalytic reactions. *Physical Review E*, 79:036112, 2009.
- [208] Y. Togashi and K. Kaneko. Transitions induced by the discreteness of molecules in a small autocatalytic system. *Physical Review Letters*, 86:2459–2462, 2001.
- [209] J. Zhao, L. Yeo, E. W. Huang, and P. W. Phillips. Thermodynamics of an exactly solvable model for superconductivity in a doped mott insulator. *Physical Review B*, 105(18):184509, 2022.
- [210] J. Zhao, P. Mai, B. Bradlyn, and P. Phillips. Failure of topological invariants in strongly correlated matter. *Physical review letters*, 131(10):106601, 2023.
- [211] J. Zhao, G. La Nave, and P. W. Phillips. Proof of a stable fixed point for strongly correlated electron matter. *Physical Review B*, 108(16):165135, 2023.
- [212] P. Mai, J. Zhao, G. Tenkila, N. A. Hackner, D. Kush, D. Pan, and P. W. Phillips. New approach to strong correlation: Twisting Hubbard into the orbital Hatanagai-Kohmoto model. *arXiv preprint arXiv:2401.08746*, 2024.
- [213] Y. Ma, J. Zhao, E. W. Huang, D. Kush, B. Bradlyn, and P. W. Phillips. Charge susceptibility and kubo response in Hatanagai-Kohmoto-related models. *Physical Review B*, 112(4):045109, 2025.
- [214] G. La Nave, J. Zhao, and P. W. Phillips. The Luttinger count is the homotopy not the physical charge: Generalized anomalies characterize non-fermi liquids. *arXiv preprint arXiv:2506.04342*, 2025.
- [215] A. Jain, K. Jensen, R. Liu, and E. Mefford. Dipole superfluid hydrodynamics. *Journal of High Energy Physics*, 2023(9):1–67, 2023.
- [216] A. Jain, K. Jensen, R. Liu, and E. Mefford. Dipole superfluid hydrodynamics. part ii. *Journal of High Energy Physics*, 2024(7):1–50, 2024.
- [217] C. Stahl, M. Qi, P. Glorioso, A. Lucas, and R. Nandkishore. Fracton superfluid hydrodynamics. *Physical Review B*, 108(14):144509, 2023.
- [218] P. Glorioso, X. Huang, J. Guo, J. F. Rodriguez-Nieva, and A. Lucas. Goldstone bosons and fluctuating hydrodynamics with dipole and momentum conservation. *Journal of High Energy Physics*, 2023(5):1–43, 2023.
- [219] Y. Yang, V. Gorelov, C. Pierleoni, D. M. Ceperley, and M. Holzmann. Electronic band gaps from quantum Monte Carlo methods. *Physical Review B*, 101(8):085115, 2020.
- [220] Y. Yang, N. Hiraoka, K. Matsuda, M. Holzmann, and D. M. Ceperley. Quantum Monte Carlo compton profiles of solid and liquid lithium. *Physical Review B*, 101(16):165125, 2020.
- [221] N. Hiraoka, Y. Yang, T. Hagiya, A. Niozu, K. Matsuda, S. Huotari, M. Holzmann, and D. Ceperley. Direct observation of the momentum distribution and renormalization factor in lithium. *Physical Review B*, 101(16):165124, 2020.

- [222] M. Holzmann, R. C. Clay III, M. A. Morales, N. M. Tubman, D. M. Ceperley, and C. Pierleoni. Theory of finite size effects for electronic quantum Monte Carlo calculations of liquids and solids. *Physical Review B*, 94(3):035126, 2016.
- [223] N. Drummond, R. Needs, A. Sorouri, and W. Foulkes. Finite-size errors in continuum quantum Monte Carlo calculations. *Physical Review B—Condensed Matter and Materials Physics*, 78(12):125106, 2008.
- [224] S. Chiesa, D. M. Ceperley, R. M. Martin, and M. Holzmann. Finite-size error in many-body simulations with long-range interactions. *Physical review letters*, 97(7):076404, 2006.
- [225] P. Gérard and E. Lenzmann. A lax pair structure for the half-wave maps equation. *Letters in Mathematical Physics*, 108(7):1635–1648, 2018.
- [226] E. Lenzmann and J. Sok. Derivation of the half-wave maps equation from Calogero–Moser spin systems. *arXiv preprint arXiv:2007.15323*, 2020.
- [227] T. Zhou and M. Stone. Solitons in a continuous classical Haldane–Shastry spin chain. *Physics Letters A*, 379(43–44):2817–2825, 2015.
- [228] E. Lenzmann. A short primer on the half-wave maps equation. pages, 1–12, 2018.
- [229] R. P. Stanley. What is enumerative combinatorics? pages, 1–63. Springer, 1986.
- [230] R. Simion. Noncrossing partitions. *Discrete Mathematics*, 217(1–3):367–409, 2000.
- [231] T. Zhou and A. Nahum. Emergent statistical mechanics of entanglement in random unitary circuits. *Physical Review B*, 99(17):174205, 2019.
- [232] B. Hsu and E. Fradkin. Dynamical stability of the quantum Lifshitz theory in $2+1$ dimensions. *Physical Review B—Condensed Matter and Materials Physics*, 87(8):085102, 2013.
- [233] E. Fradkin. Scaling of entanglement entropy at 2d quantum Lifshitz fixed points and topological fluids. *Journal of Physics A: Mathematical and Theoretical*, 42(50):504011, 2009.
- [234] D. E. Parker, R. Vasseur, and J. E. Moore. Entanglement entropy in excited states of the quantum Lifshitz model. *Journal of Physics A: Mathematical and Theoretical*, 50(25):254003, 2017.
- [235] É. Brunet. *Some aspects of the Fisher-KPP equation and the branching Brownian motion*. PhD thesis, UPMC, 2016.
- [236] S. Chatterjee and P. S. Dey. Multiple phase transitions in long-range first-passage percolation on square lattices. *Communications on Pure and Applied Mathematics*, 69(2):203–256, 2016.
- [237] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić. Weak ergodicity breaking from quantum many-body scars. *Nature Physics*, 14(7):745–749, 2018.
- [238] A. Nahum, J. Ruhman, S. Vijay, and J. Haah. Quantum entanglement growth under random unitary dynamics. *Physical Review X*, 7(3):031016, 2017.
- [239] T. Zhou and D. J. Luitz. Operator entanglement entropy of the time evolution operator in chaotic systems. *Physical Review B*, 95(9):094206, 2017.
- [240] H. P. NOYES. page, 1. North-Holland Publishing Company, 1970.
- [241] V. Efimov. Effective interaction of three resonantly interacting particles and the force range. *Phys. Rev. C*, 47:1876–1884, 1993.
- [242] Y. Castin and F. Werner. Single-particle momentum distribution of an Efimov trimer. *Phys. Rev. A*, 83:063614, 2011.
- [243] V. E. Colussi, J. P. Corson, and J. P. D’Incao. Dynamics of three-body correlations in quenched unitary bose gases. *Phys. Rev. Lett.*, 120:100401, 2018.

- [244] V. E. Colussi, B. E. van Zwol, J. P. D’Incao, and S. J. J. M. F. Kokkelmans. Bunching, clustering, and the buildup of few-body correlations in a quenched unitary bose gas. *Phys. Rev. A*, 99:043604, 2019.
- [245] C. Fefferman and C. R. Graham. *The ambient metric (AM-178)*. Princeton University Press, 2012.
- [246] C. R. Graham. Extended obstruction tensors and renormalized volume coefficients. *Advances in Mathematics*, 220(6):1956–1985, 2009.
- [247] W. Jia and M. Karydas. Obstruction tensors in Weyl geometry and holographic Weyl anomaly. *Physical Review D*, 104(12):126031, 2021.
- [248] W. Jia, M. Karydas, and R. G. Leigh. Weyl-ambient geometries. *Nuclear Physics B*, 991:116224, 2023.
- [249] M. Henningson and K. Skenderis. The holographic Weyl anomaly. *Journal of High Energy Physics*, 1998(07):023, 1998.
- [250] L. Ciambelli and R. G. Leigh. Weyl connections and their role in holography. *Physical Review D*, 101(8):086020, 2020.
- [251] M. Safronova, S. Porsev, and C. W. Clark. Ytterbium in quantum gases and atomic clocks: van der Waals interactions and blackbody shifts. *Physical review letters*, 109(23):230802, 2012.
- [252] J. Mitroy, M. S. Safronova, and C. W. Clark. Theory and applications of atomic and ionic polarizabilities. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(20):202001, 2010.
- [253] F. Le Kien, P. Schneeweiss, and A. Rauschenbeutel. Dynamical polarizability of atoms in arbitrary light fields: general theory and application to cesium. *The European Physical Journal D*, 67(5): 92, 2013.
- [254] Z.-M. Tang, Y.-M. Yu, J. Jiang, and C.-Z. Dong. Magic wavelengths for the $6s^{21}S_0-6s6p^3P_1^o$ transition in ytterbium atom. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 51(12):125002, 2018.
- [255] M. Fujita, T. Takayanagi, and E. Tonni. Aspects of AdS/BCFT. *Journal of High Energy Physics*, 2011(11):1–40, 2011.
- [256] M. Grinberg and J. Maldacena. Proper time to the black hole singularity from thermal one-point functions. *Journal of High Energy Physics*, 2021(3):1–31, 2021.
- [257] H. Geng and Y. Jiang. Microscopic origin of the entropy of single-sided black holes. *Journal of High Energy Physics*, 2025(4):1–29, 2025.
- [258] W. Z. Chua and Y. Jiang. Hartle-Hawking state and its factorization in 3d gravity. *Journal of High Energy Physics*, 2024(3):1–81, 2024.
- [259] S. Horvat and B. Dakić. Interference as an information-theoretic game. *Quantum*, 5:404, 2021.
- [260] X. Chen, Y. Zhang, A. Winter, V. O. Lorenz, and E. Chitambar. Information carried by a single particle in quantum multiple-access channels. *Physical Review A*, 109(6):062420, 2024.
- [261] J. Maisriml, S. Horvat, and B. Dakić. Acquisition of delocalized information via classical and quantum carriers. *arXiv preprint arXiv:2506.11254*, 2025.
- [262] S. Kochen and E. P. Specker. The problem of hidden variables in quantum mechanics. pages, 235–263. Springer, 2011.
- [263] R. W. Spekkens. Contextuality for preparations, transformations, and unsharp measurements. *Physical Review A—Atomic, Molecular, and Optical Physics*, 71(5):052108, 2005.
- [264] A. A. Klyachko, M. A. Can, S. Binicioğlu, and A. S. Shumovsky. Simple test for hidden variables in spin-1 systems. *Physical review letters*, 101(2):020403, 2008.

- [265] R. Kunjwal and R. W. Spekkens. From the Kochen-Specker theorem to noncontextuality inequalities without assuming determinism. *Physical review letters*, 115(11):110403, 2015.
- [266] Y. Zhang, Y. Ying, and D. Schmid. Quantifiers and witnesses for the nonclassicality of measurements and of states. *arXiv preprint arXiv:2504.02944*, 2025.
- [267] Y. Zhang, D. Schmid, Y. Ying, and R. Spekkens. Reassessing the boundary between classical and nonclassical for individual quantum processes, arxiv (2025). *arXiv preprint arXiv:2503.05884*.
- [268] V. S. Dotsenko and V. A. Fateev. Conformal algebra and multipoint correlation functions in 2d statistical models. *Nuclear Physics B*, 240(3):312–348, 1984.
- [269] R. Dijkgraaf, E. Verlinde, and H. Verlinde. $C=1$ conformal field theories on Riemann surfaces. *Communications in Mathematical Physics*, 115(4):649–690, 1988.
- [270] O. Aharony, J. Marsano, S. Minwalla, K. Papadodimas, and M. V. RAAMSDONK. The Hagedorn/deconfinement phase transition in weakly coupled large n gauge theories. pages, 161–203. World Scientific, 2004.
- [271] J. Kinney, J. Maldacena, S. Minwalla, and S. Raju. An index for 4 dimensional super conformal theories. *Communications in mathematical physics*, 275(1):209–254, 2007.
- [272] J. Yu, J. Herzog-Arbeitman, and B. A. Bernevig. Universal Wilson loop bound of quantum geometry. *Physical Review Letters*, 135(8):086401, 2025.
- [273] J. Yu, B. Lian, and S. Ryu. Wilson-loop-ideal bands and general idealization. *arXiv preprint arXiv:2509.05410*, 2025.
- [274] J. Yu, B. A. Bernevig, R. Queiroz, E. Rossi, P. Törmä, and B.-J. Yang. Quantum geometry in quantum materials. *arXiv preprint arXiv:2501.00098*, 2024.
- [275] R. Roy. Band geometry of fractional topological insulators. *Physical Review B*, 90(16):165139, 2014.
- [276] K. Yang, Y. Liu, F. Schindler, and C.-X. Liu. Engineering miniband topology via band folding in moiré superlattice materials. *Physical Review B*, 111(24):L241104, 2025.
- [277] C. Hernández-García, A. Picón, J. San Román, and L. Plaja. Attosecond extreme ultraviolet vortices from high-order harmonic generation. *Physical review letters*, 111(8):083602, 2013.
- [278] K. M. Dorney, L. Rego, N. J. Brooks, J. San Román, C.-T. Liao, J. L. Ellis, D. Zusin, C. Gentry, Q. L. Nguyen, J. M. Shaw, et al. Controlling the polarization and vortex charge of attosecond high-harmonic beams via simultaneous spin-orbit momentum conservation. *Nature photonics*, 13(2):123–130, 2019.
- [279] L. Rego, K. M. Dorney, N. J. Brooks, Q. L. Nguyen, C.-T. Liao, J. San Román, D. E. Couch, A. Liu, E. Pisanty, M. Lewenstein, et al. Generation of extreme-ultraviolet beams with time-varying orbital angular momentum. *Science*, 364(6447):eaaw9486, 2019.
- [280] N. J. Brooks, A. de las Heras, B. Wang, I. Binnie, J. Serrano, J. San Román, L. Plaja, H. C. Kapteyn, C. Hernandez-Garcia, and M. M. Murnane. Circularly polarized attosecond pulses enabled by an azimuthal phase and polarization grating. *ACS Photonics*, 12(1):495–504, 2024.
- [281] A. Hook, Y. Kahn, B. R. Safdi, and Z. Sun. Radio signals from axion dark matter conversion in neutron star magnetospheres. *Phys. Rev. Lett.*, 121:241102, 2018.
- [282] M. Leroy, M. Chianese, T. D. P. Edwards, and C. Weniger. Radio signal of axion-photon conversion in neutron stars: A ray tracing analysis. *Phys. Rev. D*, 101:123003, 2020.
- [283] S. J. Witte, D. Noordhuis, T. D. P. Edwards, and C. Weniger. Axion-photon conversion in neutron star magnetospheres: The role of the plasma in the Goldreich-Julian model. *Phys. Rev. D*, 104:103030, 2021.
- [284] H. Grote and Y. V. Stadnik. Novel signatures of dark matter in laser-interferometric gravitational-wave detectors. *Phys. Rev. Res.*, 1:033187, 2019.

- [285] L. Aiello, J. W. Richardson, S. M. Vermeulen, H. Grote, C. Hogan, O. Kwon, and C. Stoughton. Constraints on scalar field dark matter from colocated Michelson interferometers. *Phys. Rev. Lett.*, 128:121101, 2022.
- [286] S. M. Vermeulen, P. Relton, H. Grote, V. Raymond, C. Affeldt, F. Bergamin, A. Bisht, M. Brinkmann, K. Danzmann, S. Doravari, V. Kringel, J. Lough, H. Lück, M. Mehmet, N. Mukund, S. Nadji, E. Schreiber, B. Sorazu, K. A. Strain, H. Vahlbruch, M. Weinert, B. Willke, and H. Wittel. Direct limits for scalar field dark matter from a gravitational-wave detector. *Nature*, 600(7889): 424–428, 2021.
- [287] E. Hall and N. Aggarwal. Advanced LIGO, LISA, and Cosmic Explorer as dark matter transducers, 2022.
- [288] S. Morisaki, T. Fujita, Y. Michimura, H. Nakatsuka, and I. Obata. Improved sensitivity of interferometric gravitational-wave detectors to ultralight vector dark matter from the finite light-traveling time. *Phys. Rev. D*, 103:L051702, 2021.
- [289] A. Pierce, K. Riles, and Y. Zhao. Searching for dark photon dark matter with gravitational-wave detectors. *Phys. Rev. Lett.*, 121:061102, 2018.
- [290] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, et al. Constraints on dark photon dark matter using data from LIGO’s and Virgo’s third observing run. *Phys. Rev. D*, 105:063030, 2022.
- [291] A. L. Miller, P. Astone, G. Bruno, S. Clesse, S. D’Antonio, A. Depasse, F. De Lillo, S. Frasca, I. La Rosa, P. Leaci, C. Palomba, O. J. Piccinni, L. Pierini, L. Rei, and A. Tanasijczuk. Probing new light gauge bosons with gravitational-wave interferometers using an adapted semicoherent method. *Phys. Rev. D*, 103:103002, 2021.
- [292] V. Sunko, Y. Sun, M. Vranas, C. C. Homes, C. Lee, E. Donoway, Z.-C. Wang, S. Balguri, M. B. Mahendru, A. Ruiz, et al. Spin-carrier coupling induced ferromagnetism and giant resistivity peak in EuCd_2P_2 . *Physical Review B*, 107(14):144404, 2023.
- [293] E. Donoway, T. Trevisan, A. Liebman-Peláez, R. Day, K. Yamakawa, Y. Sun, J. Soh, D. Prabhakaran, A. Boothroyd, R. Fernandes, et al. Multimodal approach reveals the symmetry-breaking pathway to the broken helix in EuIn_2As_2 . *Physical Review X*, 14(3):031013, 2024.
- [294] E. J. Sie, C. H. Lui, Y.-H. Lee, L. Fu, J. Kong, and N. Gedik. Large, valley-exclusive Bloch-Siegert shift in monolayer WS_2 . *Science*, 355(6329):1066–1069, 2017.
- [295] E. J. Sie, J. W. McIver, Y.-H. Lee, L. Fu, J. Kong, and N. Gedik. Valley-selective optical Stark effect in monolayer WS_2 . *Nature Materials*, 14(3):290–294, 2015.
- [296] A. Overhauser. Observability of charge-density waves by neutron diffraction. *Physical Review B*, 3(10):3173, 1971.
- [297] M. Pospelov, A. Ritz, and M. B. Voloshin. Secluded WIMP dark matter. *Phys. Lett. B*, 662:53–61, 2008.
- [298] B. I. Shklovskii. Simple model of Coulomb disorder and screening in graphene. *Phys. Rev. B*, 76: 233411, 2007.
- [299] S. Adam, E. Hwang, V. Galitski, and S. Das Sarma. A self-consistent theory for graphene transport. *Proceedings of the National Academy of Sciences*, 104(47):18392–18397, 2007.
- [300] T. Li, P. Zhu, W. A. Benalcazar, and T. L. Hughes. Fractional disclination charge in two-dimensional C_n -symmetric topological crystalline insulators. *Phys. Rev. B*, 101:115115, 2020.
- [301] P. Zhu, K. Loehr, and T. L. Hughes. Identifying C_n -symmetric higher-order topology and fractional corner charge using entanglement spectra. *Phys. Rev. B*, 101:115140, 2020.
- [302] F. Zschocke and M. Vojta. Physical states and finite-size effects in Kitaev’s honeycomb model: Bond disorder, spin excitations, and nmr line shape. *Phys. Rev. B*, 92:014403, 2015.

- [303] P. Zhu, S. Feng, K. Wang, T. Xiang, and N. Trivedi. Emergent quantum Majorana metal from a chiral spin liquid. *Nature Communications*, 16(1):2420, 2025.
- [304] K. Wang, S. Feng, P. Zhu, R. Chi, H.-J. Liao, N. Trivedi, and T. Xiang. Fractionalization signatures in the dynamics of quantum spin liquids. *Phys. Rev. B*, 111:L100402, 2025.
- [305] S. Feng, P. Zhu, J. Knolle, and M. Knap. Transient localization from fractionalization: vanishingly small heat conductivity in gapless quantum magnets. *arXiv preprint arXiv:2509.07062*, 2025.
- [306] J.-J. Su and A. H. MacDonald. Spatially indirect exciton condensate phases in double bilayer graphene. *Physical Review B*, 95(4):045416, 2017.
- [307] P. Abbamonte and J. Fink. Collective charge excitations studied by electron energy-loss spectroscopy. *Annual Review of Condensed Matter Physics*, 16(1):465–480, 2025.
- [308] M. Mitrano, S. Johnston, Y.-J. Kim, and M. Dean. Exploring quantum materials with resonant inelastic x-ray scattering. *Physical Review X*, 14(4):040501, 2024.
- [309] J. Zhou, S. Chessa, E. Chitambar, and F. Leditzky. On the distinguishability of geometrically uniform quantum states, 2025.
- [310] C. Hirche, C. Rouzé, and D. S. França. On contraction coefficients, partial orders and approximation of capacities for quantum channels. *Quantum*, 6:862, 2022.
- [311] F. Wu, T. Lovorn, E. Tutuc, I. Martin, and A. H. MacDonald. Topological insulators in twisted transition metal dichalcogenide homobilayers. *Phys. Rev. Lett.*, 122:086402, 2019.
- [312] A. P. Reddy, F. Alsallom, Y. Zhang, T. Devakul, and L. Fu. Fractional quantum anomalous hall states in twisted bilayer MoTe_2 and WSe_2 . *Phys. Rev. B*, 108:085117, 2023.
- [313] C. Wang, X.-W. Zhang, X. Liu, Y. He, X. Xu, Y. Ran, T. Cao, and D. Xiao. Fractional Chern insulator in twisted bilayer MoTe_2 . *Phys. Rev. Lett.*, 132:036501, 2024.
- [314] J. Yu, J. Herzog-Arbeitman, M. Wang, O. Vafeek, B. A. Bernevig, and N. Regnault. Fractional Chern insulators versus nonmagnetic states in twisted bilayer MoTe_2 . *Phys. Rev. B*, 109:045147, 2024.
- [315] Y. Jia, J. Yu, J. Liu, J. Herzog-Arbeitman, Z. Qi, H. Pi, N. Regnault, H. Weng, B. A. Bernevig, and Q. Wu. Moiré fractional Chern insulators. i. first-principles calculations and continuum models of twisted bilayer MoTe_2 . *Phys. Rev. B*, 109:205121, 2024.
- [316] U. AI Security Institute. Inspect AI: framework for large language model evaluations, . URL https://github.com/UKGovernmentBEIS/inspect_ai.
- [317] Together AI. Together AI Platform. <https://www.together.ai/>, 2025. Accessed: 2025-08-24.