AutoReporter: Development and validation of an LLM-based tool for automated research reporting guideline assessment

David Chen[1,2], Patrick Li[3], Ealia Khoshkish[2]*, Seungmin Lee[2]*, Tony Ning[2]*, Umair Tahir[2]*, Henry CY Wong[4], Michael SF Lee[5], Srinivas Raman[6,7]**

[1] Princess Margaret Cancer Centre, Radiation Medicine Program, Toronto, Ontario, Canada
[2] Temerty Faculty of Medicine, University of Toronto, Toronto, Ontario, Canada
[3] Faculty of Engineering, McMaster University, Hamilton, Ontario, Canada
[4] Department of Oncology, Princess Margaret Hospital, Hong Kong, China
[5] Department of Radiation Oncology, National University Cancer Institute, National University Hospital, Singapore
[6] Department of Radiation Oncology, BC Cancer, Vancouver, BC Canada
[7] Division of Radiation Oncology, University of British Columbia, Vancouver, BC Canada

*Equal Contribution
**Corresponding Author

Correspondence:
Dr. Srinivas Raman MD MASc FRCPC
Department of Radiation Oncology, BC Cancer Vancouver, Vancouver, British Columbia
600 W 10th Ave, Vancouver, BC, Canada V5Z 4E6
Email: Srinivas.Raman@bccancer.bc.ca

Word Count: 1994

Abstract

Incomplete reporting of scientific research can compromise the cornerstone of reproducibility and peer review, necessary for evidence generation. Evaluation of the adherence of manuscript publications to established reporting guidelines remains a resource-intensive and subjective task. To address this task, we created AutoReporter, a large language model–based system that automates item-wise assessment of reporting guideline adherence. Benchmarks of reasoning and general-purpose LLMs across eight prompt-engineered and retrieval augmented generation method combinations compared to human ratings in the SPIRIT-CONSORT-TM corpus demonstrate that the zero□shot, no□retrieval prompt coupled with the reasoning LLM offered the best balance of mean classification performance (CONSORT accuracy: 90.09%; SPIRIT accuracy: 92.07%), run-time (CONSORT run-time: 617.26 seconds; SPIRIT run-time: 544.51 seconds), and cost (CONSORT cost: 0.68 USD; SPIRIT cost: 0.65 USD). Evaluation of AutoReporter with BenchReport, a novel benchmark dataset of 10 systematic reviews that assessed reporting guideline adherence of included studies demonstrated a mean weighted accuracy of 91.8□% and substantial inter-assessor agreement in 9 out of 10 reviews compared to human ratings. These results establish the feasibility of using LLMs to automate reporting guideline adherence assessments, enabling the potential for scalable quality control in scientific research reporting.

Introduction

Concerns about scientific reproducibility have underscored the importance of comprehensive research reporting [1,2]. Consensus guidelines such as CONSORT for randomized trials [3] and SPIRIT for trial protocols [4] aim to standardize methodological reporting, yet recent audits report modest post-guideline improvements and persistent omissions across clinical research guidelines [5,6] and emerging artificial intelligence applications [7]. There remains a translational research gap between theoretical endorsement and widespread adoption of reporting guidelines.

A significant barrier is the labor-intensive and subjective nature of guideline adherence review within a backdrop of rapidly expanding volumes of biomedical literature and novel extensions and versions of established reporting guidelines [8,9]. To address this issue, expert-led editorial intervention can improve reporting completeness but are difficult to scale [10]. Rule- and keyword-driven natural language processing (NLP) methods can detect CONSORT items with above 90% accuracy [11], but are limited in contextual and semantic understanding necessary to generalize across diverse research texts and guidelines. Machine learning (ML) classifiers show promise in classifying article types [12] and assessing risk of bias [13], motivating the application of modern ML-derived large language models (LLM) with near-human NLP competencies to automate reporting guideline assessment [14,15].

Domain-specific LLMs such as BioBERT[16] and PubMedBERT [17] have been tested in pilot case studies of guideline assessment but require large-scale resources and extensive, annotated training data. Zero-shot applications of general-purpose LLMs, such as the popular Generative Pre-Trained Transformer (GPT), have yielded promising results, yet

retrieval augmented prompt methods and reasoning-based LLMs remain largely untested across multiple guidelines and study designs [18–20]. We developed **AutoReporter**, a generalizable LLM system that automates reporting guideline adherence of manuscripts, delivering real time feedback for authors and scalable bulk screening of reporting completeness for journal editorial teams.

Methods

Datasets

The SPIRIT-CONSORT-TM corpus contains 100 English RCT reports and 100 English RCT protocols, each annotated with reporting adherence for 83 SPIRIT and CONSORT items [21] (Supplementary Figure 1). BenchReport is a novel benchmark dataset that aggregates 10,710 human ratings for 506 English articles included in 10 systematic reviews, spanning 10 reporting guidelines, 7 clinical domains, and 6 article types [22–29] (Supplementary Table 1, Supplementary Figure 1). The case study evaluation dataset comprises 10 systematic reviews first-authored or supervised by H.C.Y.W. (n=5) or M.S.F.L (n=5). Dataset methods described in the Datasets section of Supplement.

LLM and Prompt Method Design

We evaluated two general-purpose (gpt-4o, gpt-4o-mini) [30] and two reasoning (gpt-o1, gpt-o3-mini) [31] LLMs via the OpenAI API with default parameters. Inputs comprised retrieved article text, guideline item number, item description, and with few shot prompts, examples of reported items. Zero shot and few shot prompts that included examples from the SPIRIT CONSORT TM annotation guide were tested with four retrieval augmented generation (RAG) modes: no RAG, BM25, Contextual, and Hybrid RAG methods. The output of the LLM included the binary rating of item reporting (0: not wholly reported; 1: wholly reported) with quoted evidence, and if absent, a suggestion to address the reporting deficiency. Dataset methods described in LLM Design and Prompt Method Design sections of Supplement.

Evaluation Procedure

The template prompt structure was empirically tested on the CONSORT Train split (n = 10). Eight prompt and RAG method combinations were validated on CONSORT Validate (n = 20) using the general-purpose gpt 4o mini and reasoning gpt o3 mini LLMs. The top-performing zero-shot, no RAG retrieval prompt was tested on CONSORT Test (n = 20) and SPIRIT Test (n = 20) with the general-purpose gpt-4o and gpt 4o mini as well as reasoning o1 and o3 mini LLMs. The zero-shot, no RAG retrieval prompt with the reasoning o3-mini LLM was tested on the BenchReport and prospective we The same prompt was applied with gpt o3 mini to BenchReport and the case study evaluation dataset. Evaluation methods described in the Evaluation Procedure section of Supplement.

Outcomes

Automated assessments of reporting guideline adherence were compared against gold-standard human assessments using classification performance metrics including accuracy, precision, recall, and F1 score, the inter-assessor agreement metric Cohen's kappa[32], and operative metrics including cost and run-time. Mean classification performance and operative metrics were compared between different combinations of LLMs and prompting methods. Outcome methods described in the Outcomes section of Supplement.

Statistics

Summary measures of article□level means and associated 95 % CIs were reported. Differences between prompt methods and LLMs were compared using two□sided Wilcoxon tests with Benjamini–Hochberg correction. Statistical analyses were conducted using Python 3.8.9 and *scipy* 1.11.3. Statistical methods described in the Statistics section of Supplement.

Ethics

This study did not require institutional review board (IRB) approval as it exclusively utilized publicly available, de-identified datasets. No human subjects were directly involved.

Results

Through iterative testing of prompt engineering and RAG methods on the CONSORT Train set (n=10), we developed a final template prompt specifying the LLM's role, guideline details, RAG-based article texts, and zero- or few-shot reporting examples (Figure 1A). On the CONSORT Validate set (n=20), we evaluated eight combinations of prompt design (zero-shot, few-shot) and RAG retrieval methods (no retrieval, contextual retrieval, BM25S retrieval, hybrid retrieval) with generalist (gpt-4o-mini) and reasoning (gpt-o3-mini) LLMs (Supplementary Table 2). For the reasoning LLM, no retrieval outperformed BM25S and hybrid retrieval but was comparable to contextual retrieval (Figure 1B), whereas the generalist LLM showed no differences in accuracy (Figure 1C). Notably, zero-shot prompting was non-inferior to few-shot for both generalist (Figure 1D) or reasoning (Figure 1E) LLMs. The reasoning LLM with no retrieval had the highest cost but the lowest run time, while for the generalist LLM, contextual and hybrid retrieval were more resource-intensive than no retrieval or BM25S (Figure 1F). Overall, zero-shot with no retrieval was the best-performing engineered prompt and RAG combination based on classification performance and runtime, prompting further testing of the reasoning LLM with this prompt method as the AutoReporter tool.

Evaluation of AutoReporter on the CONSORT Validate (n=20) and Test (n=20) datasets showed no differences in accuracy, precision, or recall between the generalist and reasoning LLMs, confirming generalizability on hold-out data (Supplementary Figure 3). Zero-shot remained non-inferior to few-shot across both the CONSORT (Supplementary Figure 4A, Supplementary Table 2) and SPIRIT (Supplementary Figure 4B, Supplementary Table 3) Test

datasets, confirm our previous finding that zero-shot was non-inferior to few-shot prompting in the CONSORT Validate dataset. The reasoning o3-mini LLM achieved performance comparable to the o1 LLM and exceeded that of the generalist gpt-4o-mini and gpt-4o models in both CONSORT (Figure 2A) and SPIRIT Test evaluations (Figure 2B, Table 1). Performance and operational metrics on the CONSORT and SPIRIT Test datasets are provided in Supplementary Tables 3 and 4.

The o3-mini LLM with zero-shot, no-retrieval prompts consistently matched or surpassed PubMedBERT on both CONSORT (92.7% vs. 87.5% macro precision) and SPIRIT (92.9% vs. 84.6% macro precision, 90.8% vs. 79.3% macro recall, 91.1% vs. 81.0% macro F1) Test evaluations (Table 1). The o3-mini LLM also performed comparably to the o1 LLM (90.09% vs. 91.65% accuracy) at substantially lower cost (US$0.68 vs. US$5.72). These findings informed the development of AutoReporter, which uses the top-performing, cost-effective o3-mini LLM with zero-shot, no-retrieval prompting for guideline adherence assessment.

To confirm generalizability beyond SPIRIT-CONSORT-TM, we created the public BenchReport dataset, comprising 10 systematic reviews that assessed the adherence of included studies to 10 reporting guidelines, comprising a total of 506 article texts and 10,710 human ratings of reporting item adherence (Supplementary Table 4). AutoReporter achieved over 90% accuracy in 8 of 10 reviews and substantial agreement with human raters (Cohen's κ>0.6) in 9 of 10 reviews (Table 2), with an overall weighted mean accuracy of 91.8%, macro precision of 90.0%, and macro recall of 91.9% (Figure 3).

We expect that users of AutoReporter may first assess an article draft for reporting guideline adherence, make revisions to the draft text to address reporting deficiencies, and then run AutoReporter's assessment to confirm that indeed these revisions have wholly addressed the reporting deficiencies. To pilot this real-world use case, we sampled 10 systematic reviews with a mean PRISMA adherence of 73.10% (95% CI 62.72–83.47%) assessed by AutoReporter. In consultation with the primary or supervising author of each review (H.C.Y.W. [n=5], M.S.F.L [n=5]), we revised the article text based on the AutoReporter assessment to address item-specific reporting deficiencies. AutoReporter assessed all revised manuscripts to have 100% PRISMA adherence, confirming that the tool can effectively detect the edits made to address item-specific reporting deficiencies.

Discussion

In this study aimed to design AutoReporter, an LLM-based tool for automated research reporting guideline assessment, we observed that 1) zero-shot with no RAG retrieval achieves superior performance compared to few-shot and RAG retrieval techniques, 2) the AutoReporter tool achieves non-inferior performance compared to the SOTA PubMedBERT method while foregoing the need for computational resources and annotated datasets for LLM fine-tuning, and 3) the AutoReporter tool generalizes well to assessing adherence across diverse sets of reporting guidelines, clinical domains, and article types. In both the CONSORT and SPIRIT test sets the simple prompt outperformed more elaborate few-shot and RAG pipelines and

matched—or exceeded—the state□of□the□art PubMedBERT fine□tuning approach 21 while eliminating the need for large annotated corpora or GPU□intensive training.

The iterative design AutoReporter shows that a zero-shot, no RAG retrieval strategy paired with the reasoning gpt-o3-mini LLM delivers the optimal balance of classification performance, cost, and runtime. In both the CONSORT and SPIRIT test sets, our prompt strategy outperformed more complex few□shot and RAG techniques and matched—or exceeded—the state□of□the□art PubMedBERT fine□tuned approach [21], while foregoing the need for large annotated data or resource□intensive training. This finding aligns with reports of strong zero-shot reasoning on logical reasoning benchmarks [33] and clinical applications, such as patient trial matching [34] and diagnostic inference [35]. We speculate that the superiority of the no RAG retrieval method is due to the long context windows of modern LLMs, obviating the additional RAG procedures of chunk selection and ranking [36]. Notably, we observed that careful prompt structuring alone—role definition, context, and concise instructions—can optimize the latent reasoning in both general and reasoning models, supporting previous observations of the utility of instruction structure optimization for biomedical systematic review screening [37].

Although AutoReporter tool was non-inferior to the state-of-the-art fine-tuned PubMedBERT [21] on the SPIRIT-CONSORT-TM corpus, we caution that domain-specific fine-tuning remains a viable approach to design LLMs for specialized NLP tasks. Prior work has shown that prompt engineering can out-perform fine-tuning for general medical question-answering [38], patient diagnosis classification [39], and clinical named-entity recognition [40] from clinical notes. Determining the cost-benefit analysis of fine-tuning compared to zero-shot prompting alone remains a case□by□case decision. Applied to the BenchReport corpus, AutoReporter reliably achieved high classification performance across diverse reporting guidelines, clinical domains, and article types, underscoring the system's generalizability in real-world biomedical literature. We hypothesize that the underlying reasoning o3-mini LLM may Adapt to new contexts with minimal task-specific modification [41], enabling AutoReporter to be generalized to other guidelines EQUATOR network catalogue. AutoReporter can function as an AI co-reviewer that flags missing items before peer review, mirroring the dual-reviewer schema in systematic reviews [42]. Authors can iteratively refine manuscript drafts with explainable reporting assessments while journal editors can assess reporting adherence of submissions at scale.

Our evaluations used English□only text motivating future evaluation of multilingual competency. The zero□shot approach relies on the LLM's implicit understanding of biomedical language, but fine□tuning or RAG may still improve performance in technical fields. Although AutoReporter quotes supporting text, verifying the logical rationale behind each decision or adding self-verification through agentic workflows could increase interpretability. Finally, prospective trials embedded in editorial workflows are needed to quantify user acceptance and the real□world impact on manuscript quality.

AutoReporter demonstrates that a zero□shot reasoning LLM can demonstrate fast, low□cost, and accurate assessments of reporting guideline adherence, outperforming both

retrieval augmented and fine tuned domain-specific models. AutoReporter's strong performance across multiple guidelines and article types, without fine tuning or retrieval, offers an automated, generalizable solution for authors, reviewers, and journal editors to assess the completeness of scientific research reporting.

Data Availability:
Data generated and used in this study is available at https://github.com/davidchen0420/AutoReporter-Data/tree/main.

Code Availability:
Code generated and used in this study is available at https://github.com/patrickli53/autoreporter.

Disclosures:
There are no conflicts of interest related to this work.

Funding:
None.

Acknowledgements:
None.

References

1. Baker M. 1,500 scientists lift the lid on reproducibility. Nature 2016;533(7604):452–4.

2. Savović J, Jones HE, Altman DG, et al. Influence of Reported Study Design Characteristics on Intervention Effect Estimates From Randomized, Controlled Trials. Ann Intern Med 2012;157(6):429–38.

3. Moher D, Hopewell S, Schulz KF, et al. CONSORT 2010 Explanation and Elaboration: updated guidelines for reporting parallel group randomised trials. BMJ 2010;340(mar23 1):c869–c869.

4. Chan A-W, Tetzlaff JM, Altman DG, et al. SPIRIT 2013 statement: defining standard protocol items for clinical trials. Ann Intern Med 2013;158(3):200–7.

5. Plint AC, Moher D, Morrison A, et al. Does the CONSORT checklist improve the quality of reports of randomised controlled trials? A systematic review. Med J Aust 2006;185(5):263–7.

6. Tan ZW, Tan AC, Li T, et al. Has the reporting quality of published randomised controlled trial protocols improved since the SPIRIT statement? A methodological study. BMJ Open 2020;10(8):e038283.

7. Martindale APL, Llewellyn CD, De Visser RO, et al. Concordance of randomised controlled trials for artificial intelligence interventions with the CONSORT-AI reporting guidelines. Nat Commun 2024;15(1):1619.

8. Ghosn L, Boutron I, Ravaud P. Consolidated Standards of Reporting Trials (CONSORT) extensions covered most types of randomized controlled trials, but the potential workload for authors was high. J Clin Epidemiol 2019;113:168–75.

9. MiRoR network, Blanco D, Biggane AM, Cobo E. Are CONSORT checklists submitted by authors adequately reflecting what information is actually reported in published papers? Trials 2018;19(1):80.

10. Blanco D, Schroter S, Aldcroft A, et al. Effect of an editorial intervention to improve the completeness of reporting of randomised trials: a randomised controlled trial. BMJ Open 2020;10(5):e036799.

11. Wang F, Schilsky RL, Page D, et al. Development and Validation of a Natural Language Processing Tool to Generate the CONSORT Reporting Checklist for Randomized Clinical Trials. JAMA Netw Open 2020;3(10):e2014661.

12. Thomas J, McDonald S, Noel-Storr A, et al. Machine learning reduced workload with minimal risk of missing studies: development and evaluation of a randomized controlled trial classifier for Cochrane Reviews. J Clin Epidemiol 2021;133:140–51.

13. Marshall IJ, Kuiper J, Wallace BC. Automating Risk of Bias Assessment for Clinical Trials. IEEE J Biomed Health Inform 2015;19(4):1406–12.

14. Ayers JW, Poliak A, Dredze M, et al. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. JAMA

Intern Med 2023;183(6):589.

15. Chen D, Parsa R, Hope A, et al. Physician and Artificial Intelligence Chatbot Responses to Cancer Questions From Social Media. JAMA Oncol 2024;10(7):956.

16. Kilicoglu H, Rosemblat G, Hoang L, et al. Toward assessing clinical trial publications for reporting transparency. J Biomed Inform 2021;116:103717.

17. Jiang L, Lan M, Menke JD, Vorland CJ, Kilicoglu H. Text classification models for assessing the completeness of randomized controlled trial publications based on CONSORT reporting guidelines. Sci Rep 2024;14(1):21721.

18. Srinivasan A, Berkowitz J, Kivelson S, Friedrich N, Tatonetti N. Evaluating the Reporting Quality of 21,041 Randomized Controlled Trial Articles [Internet]. 2025 [cited 2025 Apr 16];Available from: http://medrxiv.org/lookup/doi/10.1101/2025.03.06.25323528

19. Li Z, Luo X, Yang Z, et al. RAPID: Reliable and efficient Automatic generation of submission rePorting checklists with Large language moDels [Internet]. 2025 [cited 2025 Apr 16];Available from: http://biorxiv.org/lookup/doi/10.1101/2025.02.13.638015

20. Wrightson JG, Blazey P, Moher D, Khan KM, Ardern CL. GPT for RCTs? Using AI to determine adherence to clinical trial reporting guidelines. BMJ Open 2025;15(3):e088735.

21. Jiang L, Vorland CJ, Ying X, et al. SPIRIT-CONSORT-TM: a corpus for assessing transparency of clinical trial protocol and results publications. Sci Data 2025;12(1):355.

22. Mc Cord KA, Imran M, Rice DB, et al. Reporting transparency and completeness in Trials: Paper 2 - reporting of randomised trials using registries was often inadequate and hindered the interpretation of results. J Clin Epidemiol 2022;141:175–86.

23. Suchá D, Van Hamersvelt RW, Van Den Hoven AF, De Jong PA, Verkooijen HM. Suboptimal Quality and High Risk of Bias in Diagnostic Test Accuracy Studies at Chest Radiography and CT in the Acute Setting of the COVID-19 Pandemic: A Systematic Review. Radiol Cardiothorac Imaging 2020;2(4):e200342.

24. Yuan T, Xiong J, Wang X, et al. The Quality of Methodological and Reporting in Network Meta-Analysis of Acupuncture and Moxibustion: A Cross-Sectional Survey. Evid Based Complement Alternat Med 2021;2021:1–15.

25. Adobes Martin M, Santamans Faustino S, Llario Almiñana I, Aiuto R, Rotundo R, Garcovich D. There is still room for improvement in the completeness of abstract reporting according to the PRISMA-A checklist: a cross-sectional study on systematic reviews in periodontology. BMC Med Res Methodol 2021;21(1):33.

26. Mazhar F, Hadi MA, Kow CS, Marran AMN, Merchant HA, Hasan SS. Use of hydroxychloroquine and chloroquine in COVID-19: How good is the quality of randomized controlled trials? Int J Infect Dis 2020;101:107–20.

27. Kim T-H, Kang JW, Jeon S-R, Ang L, Lee HW, Lee MS. Use of Traditional, Complementary and Integrative Medicine During the COVID-19 Pandemic: A Systematic Review and Meta-Analysis. Front Med 2022;9:884573.

28. Kim MM, Pound L, Steffensen I, Curtin GM. Reporting and methodological quality of systematic literature reviews evaluating the associations between e-cigarette use and cigarette smoking behaviors: a systematic quality review. Harm Reduct J 2021;18(1):121.

29. Besag FM, Vasey MJ, Lao KS, Chowdhury U, Stern JS. Pharmacological treatment for Tourette syndrome in children and adults: What is the quality of the evidence? A systematic review. J Psychopharmacol (Oxf) 2021;35(9):1037–61.

30. OpenAI, Hurst A, Lerer A, et al. GPT-4o System Card [Internet]. 2024 [cited 2025 Apr 16];Available from: https://arxiv.org/abs/2410.21276

31. OpenAI, :, Jaech A, et al. OpenAI o1 System Card [Internet]. 2024 [cited 2025 Apr 16];Available from: https://arxiv.org/abs/2412.16720

32. McHugh ML. Interrater reliability: the kappa statistic. Biochem Medica 2012;22(3):276–82.

33. Kojima T, Gu SS, Reid M, Matsuo Y, Iwasawa Y. Large Language Models are Zero-Shot Reasoners [Internet]. 2022 [cited 2025 Apr 16];Available from: https://arxiv.org/abs/2205.11916

34. Wornow M, Lozano A, Dash D, Jindal J, Mahaffey KW, Shah NH. Zero-Shot Clinical Trial Patient Matching with LLMs. NEJM AI [Internet] 2025 [cited 2025 Apr 16];2(1). Available from: https://ai.nejm.org/doi/10.1056/AIcs2400360

35. Cabral S, Restrepo D, Kanjee Z, et al. Clinical Reasoning of a Generative Artificial Intelligence Model Compared With Physicians. JAMA Intern Med 2024;184(5):581.

36. Pawar S, Tonmoy SMTI, Zaman SMM, Jain V, Chadha A, Das A. The What, Why, and How of Context Length Extension Techniques in Large Language Models -- A Detailed Survey [Internet]. 2024 [cited 2025 Apr 16];Available from: https://arxiv.org/abs/2401.07872

37. Cao C, Sang J, Arora R, et al. Development of Prompt Templates for Large Language Model–Driven Screening in Systematic Reviews. Ann Intern Med 2025;178(3):389–401.

38. Maharjan J, Garikipati A, Singh NP, et al. OpenMedLM: prompt engineering can out-perform fine-tuning in medical question-answering with open-source large language models. Sci Rep 2024;14(1):14156.

39. Zhang X, Talukdar N, Vemulapalli S, et al. Comparison of Prompt Engineering and Fine-Tuning Strategies in Large Language Models in the Classification of Clinical Notes. AMIA Jt Summits Transl Sci Proc AMIA Jt Summits Transl Sci 2024;2024:478–87.

40. Hu Y, Chen Q, Du J, et al. Improving large language models for clinical named entity recognition via prompt engineering. J Am Med Inform Assoc 2024;31(9):1812–20.

41. Huang J, Chang KC-C. Towards Reasoning in Large Language Models: A Survey [Internet]. 2022 [cited 2025 Apr 16];Available from: https://arxiv.org/abs/2212.10403

42. Waffenschmidt S, Knelangen M, Sieben W, Bühn S, Pieper D. Single screening versus conventional double screening for study selection in systematic reviews: a methodological systematic review. BMC Med Res Methodol 2019;19(1):132.

Figures:

Figure 1. Evaluation of prompt methods applied to reasoning and general-purpose LLMs and tested on the CONSORT Validation dataset splits. A) Diagram of study workflow for initial evaluation of prompt methods applied to LLMs using the CONSORT Train, Validation, and Test dataset splits. B) Classification accuracy of the reasoning LLM (o3-mini) using zero-shot and few-shot prompting, stratified based on RAG method. C) Classification accuracy of the general-purpose LLM (gpt-4o-mini) using zero-shot and few-shot prompting, stratified based on RAG method. D) Classification accuracy of the reasoning LLM (o3-mini) using RAG methods, stratified based on zero- or few-shot prompting method. E) Classification accuracy of the reasoning LLM (gpt-4o-mini) using RAG methods, stratified based on zero- or few-shot prompting method. F) Cost and Time of the reasoning LLM (o3-mini) using zero-shot prompting, stratified based on RAG method. F) Cost and Time of the general-purpose LLM (gpt-4o-mini) using zero-shot prompting, stratified based on RAG method.

Figure 2. Evaluation of the top-performing zero-shot, no RAG retrieval prompt method tested on the CONSORT Test and SPIRIT Test dataset splits. A) Accuracy, macro recall, and macro precision of classification performance of the zero-shot, no RAG retrieval prompt method applied to reasoning (o1, o3-mini) and general-purpose (gpt-4o, gpt-4o-mini) LLMs and tested on the CONSORT Test dataset. A) Accuracy, macro recall, and macro precision of classification performance of the zero-shot, no RAG retrieval prompt method applied to reasoning (o1, o3-mini) and general-purpose (gpt-4o, gpt-4o-mini) LLMs and tested on the SPIRIT Test dataset.

Figure 3. Mean and 95% of the macro precision and macro recall of AutoReporter, using the zero-shot, no RAG retrieval prompt method applied to the reasoning o3-mini LLM, evaluated on the BenchReport benchmark dataset.

Table 1. Classification and operative metrics of the zero-shot, no RAG retrieval prompt method applied to reasoning (o1, o3-mini) and general-purpose (gpt-4o, gpt-4o-mini) LLMs compared to the state-of-the-art PubMedBERT model using the CONSORT Test and SPIRIT Test dataset splits.

Test Set	Metric	GPT-4o	GPT-4o-mini	o1	o3-mini	SOTA PubMedBERT
CONSORT	Accuracy	84.630 (81.392 - 87.868)	80.535 (77.678 - 83.392)	91.650 (88.369 - 94.931)	90.090 (87.204 - 92.976)	N.R.
	Macro Precision	87.625 (84.491 - 90.759)	81.615 (78.854 - 84.375)	92.742 (89.419 - 96.066)	91.211 (88.633 - 93.790)	87.5
	Micro Precision	84.630 (81.392 - 87.868)	80.535 (77.678 - 83.392)	91.650 (88.369 - 94.931)	90.090 (87.204 - 92.976)	92.4
	Macro Recall	78.265 (74.411 - 82.118)	74.223 (71.409 - 77.038)	88.119 (83.887 - 92.351)	86.156 (83.156 - 89.156)	85.9
	Micro Recall	84.630 (81.392 - 87.868)	80.535 (77.678 - 83.392)	91.650 (88.369 - 94.931)	90.090 (87.204 - 92.976)	91.8
	Macro F1	79.237 (75.533 - 82.940)	74.430 (71.627 - 77.233)	89.305 (85.331 - 93.280)	87.342 (84.434 - 90.251)	85.8
	Micro F1	84.630 (81.392 - 87.868)	80.535 (77.678 - 83.392)	91.650 (88.369 - 94.931)	90.090 (87.204 - 92.976)	92.1
	Time (seconds)	405.529 (318.383 - 492.675)	349.822 (267.194 - 432.450)	868.741 (725.497 - 1011.985)	617.259 (499.517 - 735.000)	N.R.
	Cost (USD)	1.217 (1.057 - 1.376)	0.308 (0.249 - 0.368)	5.721 (4.766 - 6.676)	0.676 (0.576 - 0.775)	N.R.
SPIRIT	Accuracy	86.978 (85.023 - 88.933)	83.451 (81.325 - 85.578)	90.336 (87.913 - 92.759)	92.071 (89.892 - 94.250)	N.R.
	Macro Precision	88.597 (86.769 - 90.426)	84.296 (82.177 - 86.415)	90.437 (88.102 - 92.773)	92.853 (90.681 - 95.024)	84.6
	Micro Precision	86.978 (85.023 - 88.933)	83.451 (81.325 - 85.578)	90.336 (87.913 - 92.759)	92.071 (89.892 - 94.250)	91.7
	Macro Recall	84.930 (82.998 - 86.862)	81.343 (78.673 - 84.012)	89.083 (86.663 - 91.502)	90.839 (88.955 - 92.723)	79.3
	Micro Recall	86.978 (85.023 - 88.933)	83.451 (81.325 - 85.578)	90.336 (87.913 - 92.759)	92.071 (89.892 - 94.250)	87.1

	Macro F1	85.124 (83.263 - 86.985)	81.117 (78.690 - 83.544)	89.155 (86.648 - 91.662)	91.077 (89.033 - 93.121)	81.0
	Micro F1	86.978 (85.023 - 88.933)	83.451 (81.325 - 85.578)	90.336 (87.913 - 92.759)	92.071 (89.892 - 94.250)	89.4
	Time (seconds)	391.346 (333.110 - 449.582)	335.157 (285.157 - 385.157)	884.151 (793.774 - 974.528)	544.512 (484.523 - 604.500)	N.R.
	Cost (USD)	1.164 (0.906 - 1.422)	0.306 (0.195 - 0.417)	5.710 (4.184 - 7.235)	0.653 (0.483 - 0.823)	N.R.

Table 2. Classification and operative metrics of AutoReporter, the zero-shot, no RAG retrieval prompt method applied to the reasoning o3-mini LLM, across 10 systematic reviews in the BenchReport benchmark dataset.

Metric	Besag 2021	Chen 2025a	Chen 2025b	Kim 2021	Kim 2022	Martin 2021	Mazhar 2022	McCord 2021	Sucha 2020	Yuan 2021
Accuracy	0.949 (0.926 - 0.972)	0.952 (0.937 - 0.968)	0.919 (0.911 - 0.927)	0.924 (0.898 - 0.950)	0.878 (0.857 - 0.899)	0.935 (0.926 - 0.944)	0.947 (0.917 - 0.978)	0.821 (0.800 - 0.842)	0.979 (0.966 - 0.993)	0.915 (0.891 - 0.939)
Macro Precision	0.935 (0.899 - 0.972)	0.782 (0.667 - 0.896)	0.839 (0.818 - 0.861)	0.931 (0.903 - 0.959)	0.859 (0.834 - 0.885)	0.930 (0.920 - 0.940)	0.942 (0.907 - 0.977)	0.820 (0.795 - 0.844)	0.977 (0.962 - 0.992)	0.907 (0.875 - 0.939)
Micro Precision	0.949 (0.926 - 0.972)	0.952 (0.937 - 0.968)	0.919 (0.911 - 0.927)	0.924 (0.898 - 0.950)	0.878 (0.857 - 0.899)	0.935 (0.926 - 0.944)	0.947 (0.917 - 0.978)	0.821 (0.800 - 0.842)	0.979 (0.966 - 0.993)	0.915 (0.891 - 0.939)
Macro Recall	0.955 (0.938 - 0.972)	0.890 (0.765 - 1.014)	0.935 (0.925 - 0.945)	0.900 (0.871 - 0.928)	0.896 (0.875 - 0.916)	0.941 (0.933 - 0.950)	0.946 (0.910 - 0.981)	0.796 (0.773 - 0.819)	0.982 (0.970 - 0.993)	0.874 (0.835 - 0.914)
Micro Recall	0.949 (0.926 - 0.972)	0.952 (0.937 - 0.968)	0.919 (0.911 - 0.927)	0.924 (0.898 - 0.950)	0.878 (0.857 - 0.899)	0.935 (0.926 - 0.944)	0.947 (0.917 - 0.978)	0.821 (0.800 - 0.842)	0.979 (0.966 - 0.993)	0.915 (0.891 - 0.939)
Macro F1	0.938 (0.908 - 0.968)	0.809 (0.693 - 0.925)	0.869 (0.851 - 0.887)	0.905 (0.876 - 0.933)	0.863 (0.839 - 0.888)	0.928 (0.918 - 0.938)	0.941 (0.906 - 0.976)	0.791 (0.769 - 0.813)	0.978 (0.965 - 0.992)	0.886 (0.850 - 0.921)
Micro F1	0.949 (0.926 - 0.972)	0.952 (0.937 - 0.968)	0.919 (0.911 - 0.927)	0.924 (0.898 - 0.950)	0.878 (0.857 - 0.899)	0.935 (0.926 - 0.944)	0.947 (0.917 - 0.978)	0.821 (0.800 - 0.842)	0.979 (0.966 - 0.993)	0.915 (0.891 - 0.939)
Cohen's Kappa	0.878 (0.819 - 0.936)	0.626 (0.401 - 0.852)	0.741 (0.706 - 0.775)	0.812 (0.756 - 0.868)	0.731 (0.684 - 0.779)	0.861 (0.842 - 0.880)	0.883 (0.813 - 0.953)	0.590 (0.546 - 0.633)	0.957 (0.930 - 0.984)	0.773 (0.702 - 0.843)
Time (seconds)	235.723 (208.403 - 263.044)	596.373 (549.995 - 642.750)	332.939 (314.704 - 351.174)	161.617 (138.127 - 185.107)	82.747 (78.929 - 86.565)	62.785 (60.853 - 64.716)	171.716 (147.568 - 195.865)	125.084 (119.118 - 131.051)	186.723 (166.927 - 206.519)	234.863 (208.154 - 261.572)
Cost (USD)	0.140 (0.137 - 0.144)	0.332 (0.323 - 0.341)	0.505 (0.505 - 0.506)	0.257 (0.246 - 0.268)	0.073 (0.073 - 0.073)	0.020 (0.020 - 0.020)	0.188 (0.187 - 0.189)	0.150 (0.147 - 0.153)	0.179 (0.178 - 0.180)	0.257 (0.256 - 0.258)

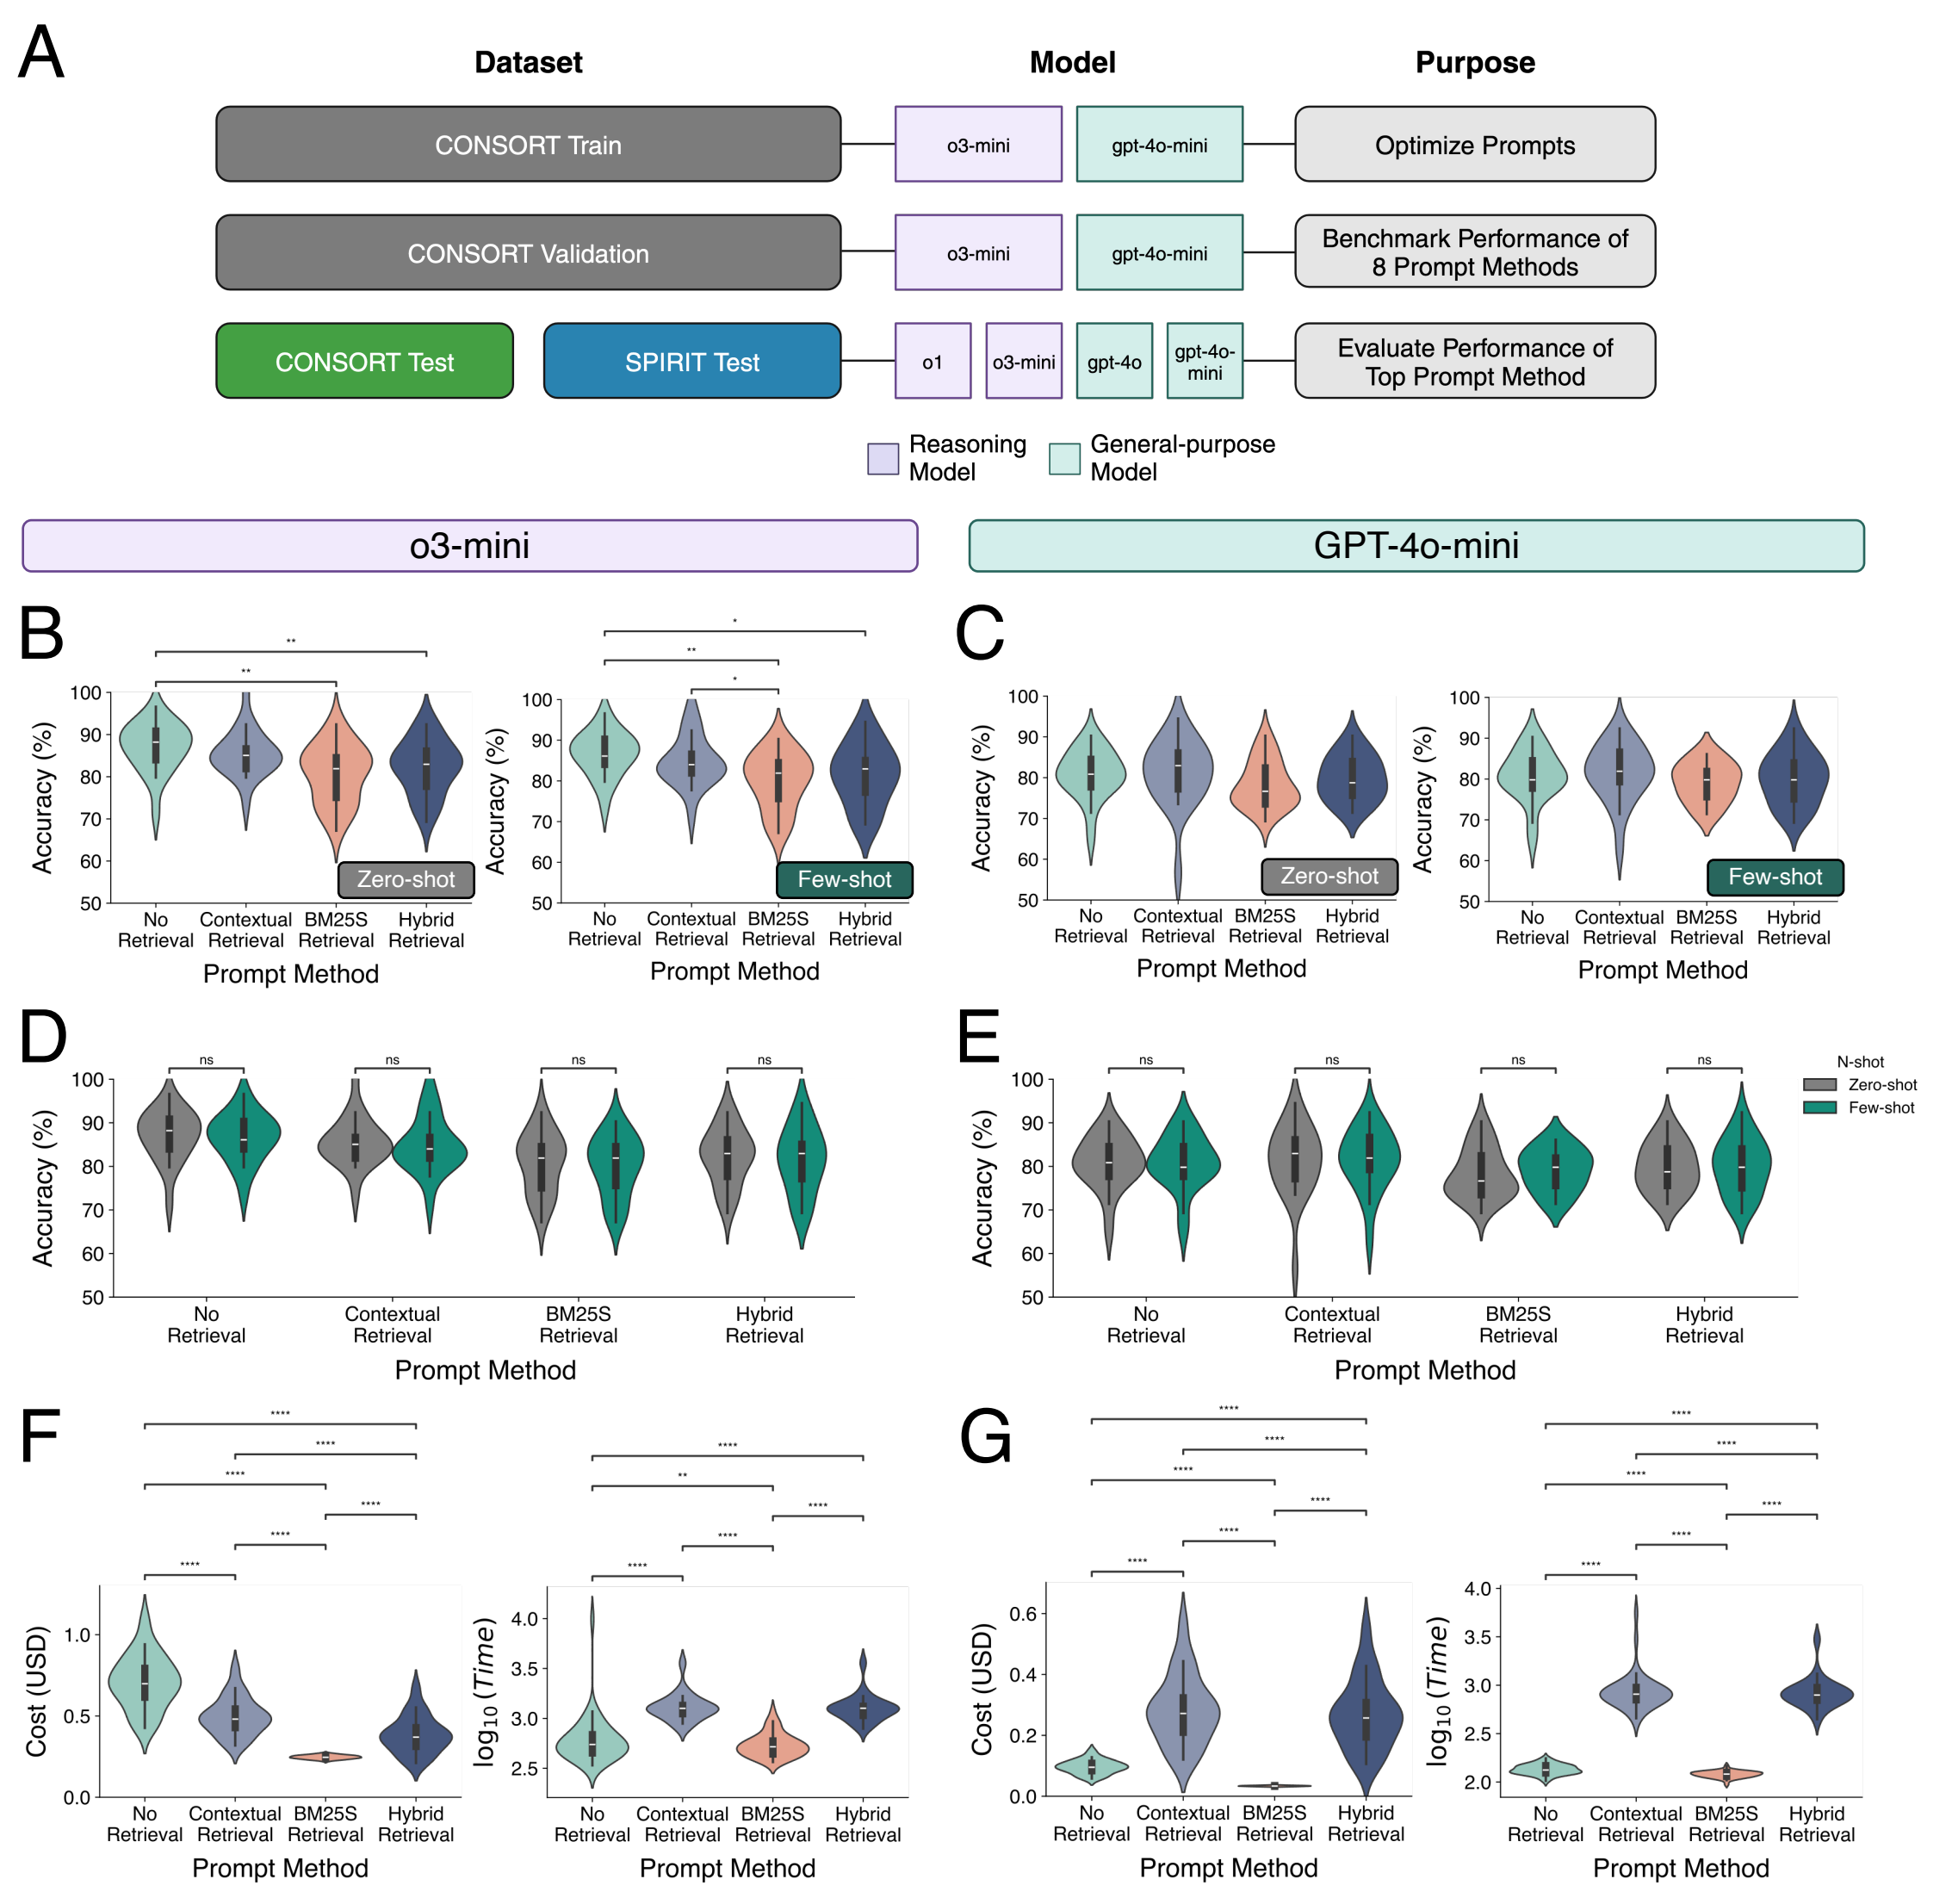
A
Dataset
Model
Purpose
CONSORT Train
o3-mini
gpt-4o-mini
Optimize Prompts
CONSORT Validation
o3-mini
gpt-4o-mini
Benchmark Performance of 8 Prompt Methods
CONSORT Test
SPIRIT Test
o1
o3-mini
gpt-4o
gpt-4o-mini
Evaluate Performance of Top Prompt Method
Reasoning Model
General-purpose Model
o3-mini
GPT-4o-mini
B
C
D
E
F
G
Accuracy (%)
Cost (USD)
log10 (Time)
Zero-shot
Few-shot
No Retrieval
Contextual Retrieval
BM25S Retrieval
Hybrid Retrieval
Prompt Method
ns
N-shot
Zero-shot
Few-shot

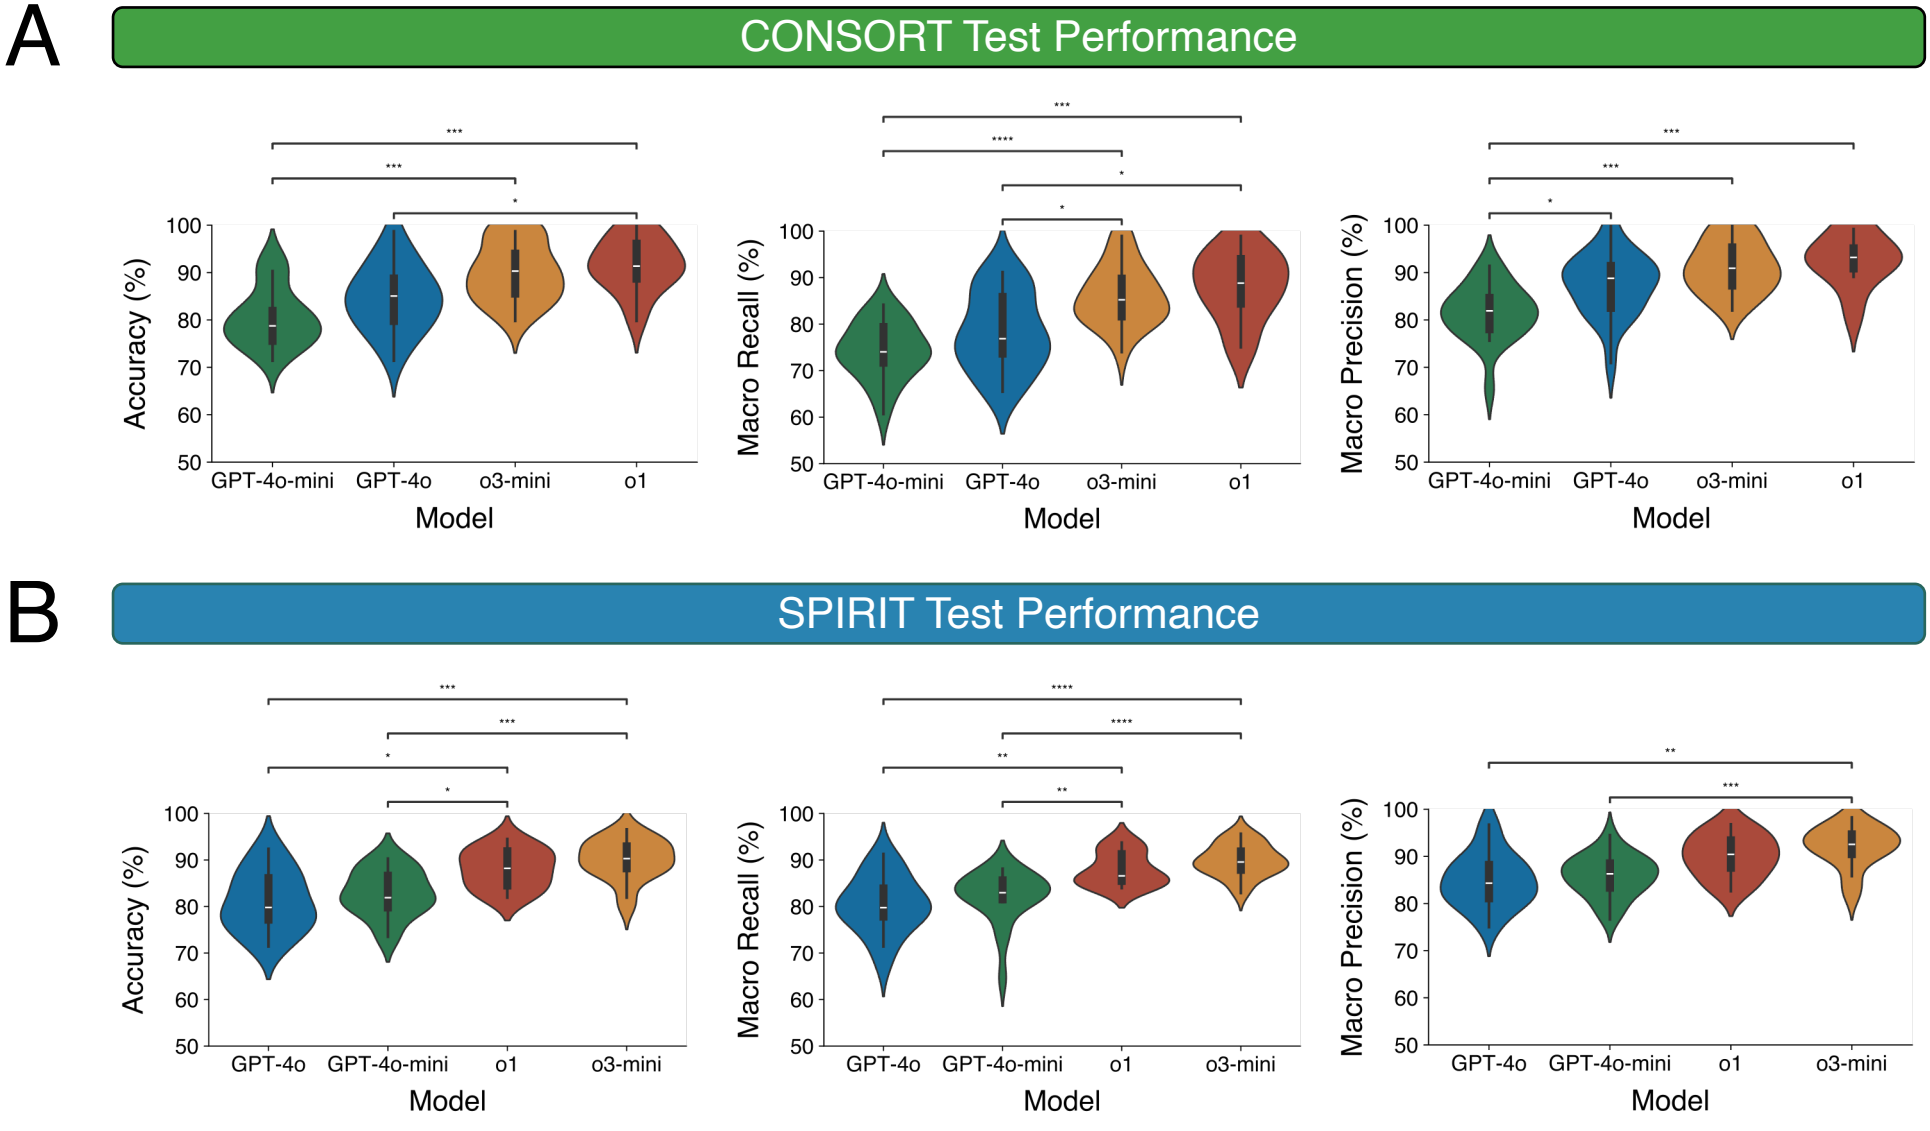
A
CONSORT Test Performance
Accuracy (%)
Macro Recall (%)
Macro Precision (%)
GPT-4o-mini
GPT-4o
o3-mini
o1
Model
B
SPIRIT Test Performance
Accuracy (%)
Macro Recall (%)
Macro Precision (%)
GPT-4o
GPT-4o-mini
o1
o3-mini
Model

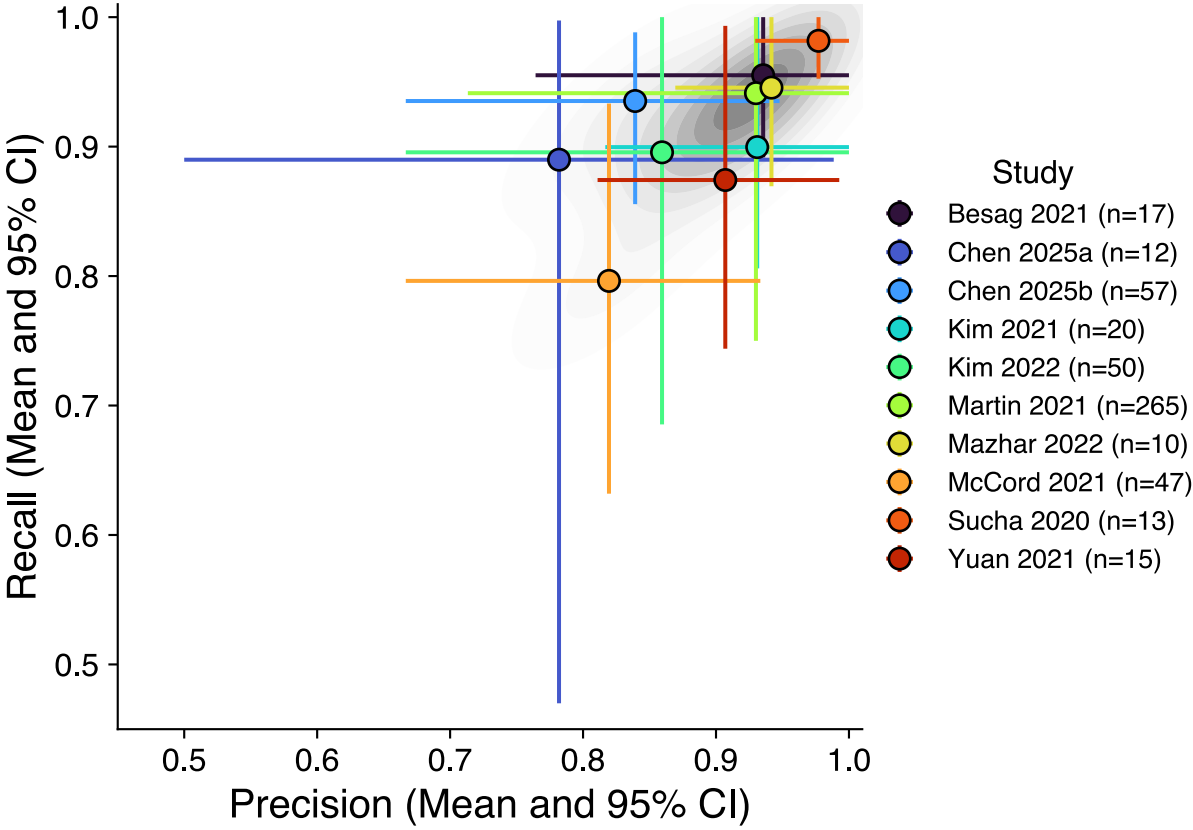
Recall (Mean and 95% CI)
1.0
0.9
0.8
0.7
0.6
0.5
Precision (Mean and 95% CI)
0.5
0.6
0.7
0.8
0.9
1.0
Study
Besag 2021 (n=17)
Chen 2025a (n=12)
Chen 2025b (n=57)
Kim 2021 (n=20)
Kim 2022 (n=50)
Martin 2021 (n=265)
Mazhar 2022 (n=10)
McCord 2021 (n=47)
Sucha 2020 (n=13)
Yuan 2021 (n=15)