



HAL
open science

CareMedEval dataset: Evaluating Critical Appraisal and Reasoning in the Biomedical Field

Doria Bonzi, Alexandre Guiggi, Frédéric Béchet, Carlos Ramisch, Benoit Favre

► To cite this version:

Doria Bonzi, Alexandre Guiggi, Frédéric Béchet, Carlos Ramisch, Benoit Favre. CareMedEval dataset: Evaluating Critical Appraisal and Reasoning in the Biomedical Field. 2025. <hal-05376617>

HAL Id: hal-05376617

<https://hal.science/hal-05376617v1>

Preprint submitted on 21 Nov 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

CareMedEval dataset: Evaluating Critical Appraisal and Reasoning in the Biomedical Field

Doria Bonzi¹, Alexandre Guiggi², Frédéric Béchet³, Carlos Ramisch³, Benoit Favre^{3,4}

¹University of Lorraine, LORIA, France

²University Grenoble-Alpes, France

³Aix-Marseille University, LIS, France

⁴CNRS, Grenoble INP, LIG, France

doria.bonzi@loria.fr, alexandre.guiggi@gmail.com, {frederic.bechet, carlos.ramisch, benoit.favre}@lis-lab.fr

Abstract

Critical appraisal of scientific literature is an essential skill in the biomedical field. While large language models (LLMs) can offer promising support in this task, their reliability remains limited, particularly for critical reasoning in specialized domains. We introduce CareMedEval, an original dataset designed to evaluate LLMs on biomedical critical appraisal and reasoning tasks. Derived from authentic exams taken by French medical students, the dataset contains 534 questions based on 37 scientific articles. Unlike existing benchmarks, CareMedEval explicitly evaluates critical reading and reasoning grounded in scientific papers. Benchmarking state-of-the-art generalist and biomedical-specialized LLMs under various context conditions reveals the difficulty of the task: open and commercial models fail to exceed an Exact Match Rate of 0.5 even though generating intermediate reasoning tokens considerably improves the results. Yet, models remain challenged especially on questions about study limitations and statistical analysis. CareMedEval provides a challenging benchmark for grounded reasoning, exposing current LLM limitations and paving the way for future development of automated support for critical appraisal.

Keywords: critical appraisal, reasoning, evaluation, medical, domain specific dataset, LLM

1. Introduction

Medical professionals must engage in continuous learning to stay up to date with evolving medical knowledge. Even though they can gather knowledge from trusted sources such as Cochrane, they also engage with latest research published in scientific papers often available before peer review in dedicated archives.

Critically appraising scientific publications is a complex cognitive task, even for trained physicians. As highlighted by previous studies, proper interpretation of biomedical literature requires not only familiarity with medical content, but also awareness of methodology and statistics (du Prel et al., 2009). Moreover, studies have shown that medical research can suffer from major methodological flaws, raising concerns about the reliability and overall trustworthiness of scientific evidence (Ioannidis, 2005; Begley and Ellis, 2012). These issues often relate to study design quality and various forms of bias (Chalmers and Glasziou, 2009; Dickersin, 1987). These challenges show why teaching and assessing critical appraisal skills remains a major and ongoing issue in the biomedical field.

Natural language processing (NLP), and in particular large language models (LLMs), represents a promising technological solution for supporting this lifelong learning process. Recent development of LLMs has put focus on their "reasoning capabilities", opening new possibilities for supporting medical professionals in the critical reading, analysis, and synthesis of scientific literature. It is important

to evaluate the reliability of such technologies.

Several evaluation benchmarks already exist in the biomedical domain, but these resources do not explicitly target the evaluation of research methodology or a system's ability to identify limitations and biases in a study. As a result, they are not well-suited for measuring the specific skills involved in critical appraisal.

Despite the growing interest in applying LLMs to these tasks, there remain significant challenges related to hallucinations, bias, and keeping scientific accuracy intact (Wang et al., 2024b; Yun et al., 2023; Meng et al., 2024). Recent work has explored long-context processing in LLMs (Nelson et al., 2024; Li et al., 2023) and their application to the medical field (Bazoge et al., 2024), including retrieval-augmented generation (RAG) systems, where external knowledge sources are integrated into the model's reasoning process (Liu et al., 2025). In these studies, RAG methods make biomedical QA more accurate and robust, yet they mainly support information access and synthesis. The task of critically assessing study design and validity is not well represented in benchmarks, especially when grounded in a given study described in a scientific article.

In this work, we introduce (1) a dataset for evaluating critical appraisal of medical studies described in scientific articles¹, derived from multiple choice question-answers (MCQA) medical education exams in French. With this dataset release, we present (2) a comprehensive evaluation of diverse state-of-the-art language models on this challeng-

ing task, providing baseline performance results and insights into model capabilities and limitations.

Unlike existing biomedical question answering datasets which are typically not grounded in research articles, our resource is directly contextualized with authentic scientific publications, aiming to evaluate the information gathering and reasoning capabilities of models. We argue that this dataset is complementary both to factual and RAG-oriented medical benchmarks, by addressing the underexplored dimension of critical appraisal in the medical field. The dataset developed in this work could also support future technologies beyond LLMs, providing a foundation for the development of tools aimed at enhancing medical reasoning and evidence-based decision-making.

2. Related Work

Systematic reviewing is a time-consuming process that has recently attracted interest for automation (Nikiforovskaya et al., 2020, Han et al., 2024). However, it has been observed that automation still faces limitations, with human reviewers outperforming current automated methods on such tasks (Yuan et al., 2022). The main goal of our study is to reflect the challenge of critical appraisal of medical articles, providing a benchmark for evaluating how well models can support both retrieval and critical reasoning in biomedical contexts.

To our knowledge, no existing dataset specifically targets the task of critical appraisal and analysis of scientific articles. However, similar resources for the general evaluation of biomedical NLP models are available: for instance, PubMedQA (Jin et al., 2019) is an English-language dataset where questions are derived from the abstracts of biomedical research articles; MedQA (Jin et al., 2020) offers an open-domain QA benchmark in English, simplified Chinese, and traditional Chinese, based on medical textbooks; SciDQA (Singh et al., 2024) provides a collection of questions grounded in full-length scientific articles with figures and images. However, none of these datasets explicitly focus on the critical evaluation of research methodology. They primarily assess factual comprehension or domain knowledge, rather than the ability to review a scientific article in terms of study design, methodology or limitations. In contrast, our dataset is specifically designed to capture and evaluate these skills through multiple-choice questions grounded in medical literature.

While the French biomedical NLP landscape still lacks a dataset that aims to explore critical appraisal of scientific articles, other resources exist for adjacent tasks: MedFrenchmark (Quercia et al., 2024)

and FrenchMedMCQA (Labrak et al., 2023) focus on MCQA in the medical domain, though without contextual grounding in articles; CAS (Grabar et al., 2018) supports information extraction tasks; and QUAERO (Névéol et al., 2014) provides annotations for named entity recognition. More broadly, DrBenchmark (Labrak et al., 2024) consolidates 20 different biomedical tasks in French, offering a comprehensive evaluation suite for LLMs.

Alongside, recent efforts have proposed RAG frameworks tailored for tasks in the biomedical field. For example, Li et al. (2024b) provide a systematic evaluation of RAG-based approaches across various biomedical applications. Studies (He et al., 2025) have shown that RAG systems significantly outperform standard LLMs on tasks such as information extraction and question answering. BiomedRAG (Li et al., 2024a) simplifies the integration of retrieved knowledge by incorporating relevant passages into LLM inputs, while BioRAG (Wang et al., 2024a) combines a scientific corpora with domain-specific embedding and hierarchical knowledge structures to improve biological question reasoning. While our approach does not employ RAG, it complements this line of research by addressing the critical appraisal of scientific articles, an aspect that remains underexplored in existing resources.

These resources reflect growing interest in evaluating LLMs in the biomedical domain, but they do not yet address the evaluation of critical reading skills, grounded in scientific literature, which our dataset aims to address.

3. Dataset overview

CareMedEval (**C**ritical **A**ppraisal and **RE**asoning **M**edical **E**valuation) is a French dataset focused on evaluating critical appraisal skills in the medical field for scientific articles, as practiced in French medical education.

This dataset is composed of 534 questions taken from *Lecture Critique d'Articles* exams (LCA, *Critical appraisal of research articles*), taken by sixth-year medical students in France. During the LCA exams, which last three hours, students are asked to answer a series of multiple-choice questions grounded in critical reading of a given scientific article. Students are required to critically analyze and interpret scientific biomedical articles, most of which are clinical studies published in peer-reviewed journals. These articles include observational studies like case-control or interventional studies such as randomized clinical trials, and cover a broad range of medical specialties like epidemiology, biostatistics or public health. Students are expected to demonstrate critical reasoning skills by classifying study types and methodological approaches, understanding the implications for clinical

¹Dataset and code available at <https://github.com/bonzid/CareMedEval>.

Sample question with context and answer choices taken from the dataset

Article, available in PDF or plain text:

The NEW ENGLAND JOURNAL of MEDICINE

ORIGINAL ARTICLE

A Randomized Trial of Convalescent Plasma in Covid-19 Severe Pneumonia

V.A. Simonovich, L.D. Burgos Pratz, P. Scibona, M.V. Beruto, M.G. Vallone, C. Vázquez, N. Savoy, D.H. Giunta, L.G. Pérez, M.-L. Sánchez, A.V. Gamarnik, D.S. Ojeda, D.M. Santoro, P.J. Camino, S. Antelo, K. Rainero, G.P. Vidiella, E.A. Miyazaki, W. Cornistein, O.A. Trabadelo, F.M. Ross, M. Spotti, G. Funtowicz, W.E. Scordo, M.H. Losso, I. Ferniot, P.E. Pardo, E. Rodriguez, P. Rucci, J. Pasquali, N.A. Fuentes, M. Esperatti, G.A. Speroni, E.C. Nannini, A. Matteaccio, H.G. Michelangelo, D. Follmann, H.C. Lane, and W.H. Belloso, for the PlasmAr Study Group*

ABSTRACT

BACKGROUND
Convalescent plasma is frequently administered to patients with Covid-19 and has been reported, largely on the basis of observational data, to improve clinical outcomes. Minimal data are available from adequately powered randomized, controlled trials.

METHODS
We randomly assigned hospitalized adult patients with severe Covid-19 pneumonia in a 2:1 ratio to receive convalescent plasma or placebo. The primary outcome was

You are a physician capable of rigorously interpreting data from medical studies. Based on the given article, answer the following multiple-choice question. Provide only the letter(s) corresponding to the correct answer(s) among: A, B, C, D, E. Your answer must strictly follow the exact format: one or more letters, separated by commas (e.g., B, E).

Question:

What is the main limitation of this study?

Answers:

- The 2:1 randomization
- The choice of the primary endpoint
- The number of subjects included
- The amount of antibodies contained in the plasma
- The fact that the study was conducted in Argentina

Figure 1: Example from the dataset showing an excerpt from a scientific article, the given instruction prompt with a corresponding question, and answer choices.

cal practice, recognizing potential biases or study limitations, and evaluating statistical evidence supporting the study conclusions. A sample of our dataset showing an excerpt from a scientific article, the given instruction prompt with a corresponding question, and answer choices is shown in Figure 1.

3.1. Dataset collection

Source data This dataset is built from two main sources:

- **Epreuves Classantes Nationales (ECN) website²**, where we can find the official national LCA exams,
- **Collège National des Enseignants de Thérapeutique (CNET) website³**, which publishes mock LCA exams reviewed and approved by an educational committee to ensure alignment with real exams.

The articles used in these exams are publicly available genuine scientific papers. We used these two websites as they provide official and mock exams, freely accessible and suitable for corpus construction. The ECN website offers real past exams, adding authenticity and relevance to our dataset.

²<https://www.cng.sante.fr/candidats/internats/concours-medicaux/etudiants/epreuves-classantes-nationales-ecn>

³<https://therap.fr/lca/>

As for the CNET website, it features training exams with professional corrections and commentary, which is particularly valuable for understanding the reasoning behind correct answers. To our knowledge, there are no other free and easily accessible online sources offering similar LCA exam content.

Languages The scientific articles are in English. The questions, answers, and justifications are in French. We also provide a preliminary English translation of the French parts of the dataset, generated using Gemini 2.5 Flash.

Annotations and labels Each question in the CareMedEval dataset was manually annotated with one or more labels, reflecting the cognitive and analytical skills required to answer it. Question labeling was conducted by one expert annotator with a background in general practice, using reference textbooks and their medical expertise. The labels are defined in Table 1. Since questions often target multiple dimensions of skills, they can be assigned more than one label. These labels can help determine whether certain categories of questions are more challenging than others for models, depending on the types of skills they require.

For data sourced from the ECN website, we manually corrected the exam questions with the help of a general practitioner, as the official answer keys were not publicly available for these past exams.

For data from CNET website, we also collected the correct answers and justifications provided by

Label	Description	Skills required	Support
design	Identification of study design	Information retrieval	105
statistics	Understanding and interpretation of statistics	General knowledge, Information retrieval	239
methodology	Knowledge of scientific methodology	General conceptual understanding	219
limitations	Critical review of biases and limitations	Contextual reasoning	132
applicability	Clinical relevance and applicability	Contextual reasoning	115

Table 1: Labels by type of reasoning involved in the critical appraisal of biomedical articles. Each question in the dataset was annotated with one or more labels reflecting the cognitive and analytical skills required to answer it.

medical professionals for a subset of questions. These justifications explain why some answers are correct or false, offering valuable insight into clinical reasoning and critical appraisal.

3.2. Dataset structure

CareMedEval is composed of 534 questions in JSON format and 37 articles available in PDF. Statistics of this dataset are shown in Table 2. On average, each question contains 15.6 tokens and has 2.60 correct answers. Most questions have multiple correct answers: about 29% have two or three correct options, 20% have four, 19% only have one, and a small portion (around 3%) have five correct answers.

Statistic	Value
Questions (total)	534
With justifications	204
Vocabulary size	1,273 words
Question length (avg.)	15.6 tokens
Correct answers (avg.)	2.6
Articles (total)	37
Questions per article (avg.)	14.4
min / max	8 / 16
Article length (avg.)	5,675 tokens
min / max	2,747 / 8,332
Article length (PDF, avg.)	10 pages
Abstract length (avg.)	1,019 tokens
min / max	276 / 1,832

Table 2: CareMedEval dataset statistics.

Each question in the dataset is represented as a JSON object with the fields described in Table 3.

Each question in the dataset is linked to a scientific article through the `id_article` field, which serves as a reference key to the corresponding article file. There are on average 14.4 questions per article. The minimum is 8 and the maximum is 16. Articles are available in plain text (.txt) format for easier processing, and the original PDF versions are also included. The plain text files were generated using the PyMuPDF library⁴, which extracts raw textual content from each page, without figures. We manually reviewed and corrected the formatting issues in the resulting files to ensure readability and to preserve the original structure of the articles as much as possible. Using the Byte-Pair Encoding tokenizer from the

tiktoken⁵ library, we found that scientific articles in our dataset contain between 2,747 and 8,332 tokens, with an average length of 5,675 tokens per article. For the abstracts alone, also included in our dataset and matched with their corresponding article IDs, the number of tokens ranges from 276 to 1,832, with an average length of 1,019 tokens.

For the PDF versions, articles are on average approximately 10 pages long, with around 5,400 words and 36,000 characters. On average, each article includes about 3.3 figures, for an estimated total of 123 figures across the entire dataset (approximations computed using the PyMuPDF library). The publication years range from 2007 to 2023. In addition of the plain text and PDF versions, each question includes a direct `article_link` field pointing to the online HTML version of the article.

4. Benchmark

To systematically assess the critical appraisal abilities of LLMs, we introduce a dedicated benchmark built upon our dataset. It provides a structured evaluation framework that integrates multiple metrics, controlled scenarios, and a representative set of models to enable comprehensive and reproducible analysis.

4.1. Metrics

Our benchmark relies on four different metrics chosen to reflect the multiple-choice question-answer nature of the dataset:

- **Exact Match Ratio** (EMR) measures the proportion of questions for which the predicted set of answers exactly matches the gold standard.
- **F1-score** is the harmonic mean of precision and recall, computed between predicted and gold answer sets.
- **Hamming score** evaluates the proportion of correctly predicted labels (answer options) over the total number of possible labels, averaged over all questions.

⁴<https://github.com/pymupdf/PyMuPDF>

⁵<https://github.com/openai/tiktoken>

Field	Description
id	A unique identifier for the question
id_article	Internal article ID
source_exam	URL of the exam or online resource
date_exam	Date of the exam
article_link	URL of the article
article_date	Article publication date
question	Question as presented in the exam
answers	Dictionary mapping each option label (A to E) to its full answer text
correct_answers	List of correct answers labels
essential_answers	List of essential answers for LCA grading (23)
unacceptable_answers	List of inadmissible answers for LCA grading (19)
labels	Labels describing skills or knowledge required (see Table 1)
justification	Expert-written explanation of correct and incorrect answers (204)
nb_correct_answers	Number of correct answer options for the question

Table 3: Field descriptions for the CareMedEval dataset. Each entry corresponds to a multiple-choice question associated with a biomedical research article. Some fields are specific to certain subsets and are designed to support fine-grained evaluation, reasoning analysis, and expert-based interpretation. Subset size is specified in parenthesis.

- **LCA score** is a custom metric inspired by the grading system used in the original LCA exam from which our dataset is derived. It reflects exam-style grading: for each question, a perfect match yields 1 point, one mismatch 0.5, two mismatches 0.25, and more than two mismatches or no response 0. In addition, there are two other constraints to the LCA grading system: if a required (essential) answer is missing, the score is automatically 0, regardless of other matches. If an unacceptable answer (one that should never be selected) is included, the score is 0. The final LCA score is averaged over all questions.

In France, the LCA exam is part of the "*Epreuves de Connaissances*" ("Knowledge Exams"), which account for 60% of the final score in the national medical competition. Results for the LCA exam and for these knowledge exams in general are not publicly available. However, we know that students must achieve a minimum score of 14/20 (70%) to advance to the next stage of the competition⁶, which can be used as a reference point for assessing models' success on the task.

These metrics allow us to distinguish between full correctness (EMR), partial correctness (F1, Hamming), and real-world grading fairness (LCA score).

4.2. Evaluation scenarios and prompts

We designed multiple evaluation scenarios in which the model was given an instruction-style prompt in French specifying the expected output format, a question, and the answer choices. Depending on the setting, the full article was provided, or only the abstract was included, or no article content was given. This pipeline is shown in Figure 2.

⁶<https://www.cng.sante.fr/epreuves-dematerialisees-nationales-edn>

We used a role-oriented instruction prompt, framing the model as a medical professional. While the impact of such role framing remains debated (Zheng et al., 2024), we kept it consistent across evaluation scenarios.

4.3. Models

To enable a comprehensive evaluation on this dataset, we selected models of varying sizes (from 8B to 120B parameters), architectures, domain specialization (general-purpose vs. biomedical-tuned), reasoning token generation capabilities, allowing us to assess the effect of the benchmark on different dimensions: Qwen3-8B/32B (Yang et al., 2025), II-Medical-8B (Intelligent Internet, 2025), Gemma3-27B-text-IT (Google, 2025), MedGemma-27B-text-IT (Google, 2025). We also evaluated GPT-4.1 (OpenAI et al., 2024), GPT-4o-mini (OpenAI et al., 2024) and GPT-OSS-20B/120B (OpenAI, 2025) to assess how frontier and derived models perform compared to smaller or domain-specific alternatives.

5. Results and analysis

We conducted a series of experiments on the CareMedEval dataset to evaluate the critical appraisal skills of a diverse pool of LLMs. While this dataset can support multiple experimental setups, we focus here on a MCQA task based solely on the textual content of the articles.

All experiments were run on an cluster equipped with NVIDIA L40-48GB and A100-80GB GPUs. Inference was performed using the vLLM (Kwon et al., 2023) and Ollama⁷ engines.

In this section, we showcase some of the findings, while the complete set of results is available

⁷<https://ollama.com/>

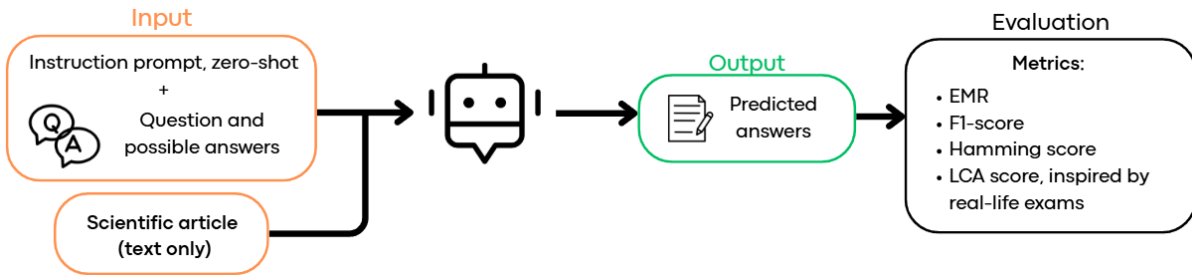


Figure 2: Overview of the model evaluation pipeline of the CareMedEval benchmark. The input consists of a zero-shot instruction prompt containing a question and possible answer choices, along with article (plain text only in our experiment setting). The model generates predicted answers, which are then evaluated using a set of quantitative metrics to assess performance.

in the [accompanying material](#). Overall, GPT-4.1 demonstrates the highest performance across all evaluation scenarios, with Qwen3-32B consistently ranking as the second best. Notably, only four models surpass an EMR of 0.25 and none surpasses a LCA score of 0.70, which would be the minimal mark for the exam, highlighting the difficulty of the task. For comparison, GPT-4.1 reaches an EMR of 0.79 on the FrenchMedMCQA dataset (Labrak et al., 2023). The difficulty of French medical MCQs from licensing examinations was highlighted in a previous study (Alfertschofer et al., 2023), where the "multiple correct answers" format appeared to be a factor contributing to models' poor performance on the task. This observation is consistent with our analysis of the subset presented in [Influence of article access on performance](#): models tend to perform better when the question explicitly requires a single correct answer. Interestingly, models specialized for the biomedical domain do not consistently outperform generalist models, exhibiting comparable performance at best.

5.1. Comparison of models performance

5.1.1. Generalist vs. specialized

The goal is to assess whether biomedical-pretrained or fine-tuned models offer a concrete advantage on medical tasks compared to the models they have been specialized from, under our hypothesis that specialized models are expected to perform better on a domain-specific critical reasoning task. Results presented in [Table 4](#) show comparable EMR scores between some specialized and generalist models (*e.g.*, Medgemma vs. Gemma3). In several cases, generalist models even outperform their specialized counterparts, such as Qwen3-8B vs. II-Medical-8B. To evaluate whether EMR differences between generalist and specialized models were significant, we applied McNemar's test, which compares paired outcomes by counting where one model is correct and

the other is not. Despite noticeable differences in raw performance, the test shows that, in most cases, these gaps are not statistically significant ($p \geq 0.05$). Significant differences were only observed for Qwen3-8B and II-Medical-8B. Consequently, our benchmark does not provide sufficient evidence to confirm our hypothesis. However, in line with prior works (Labrak et al., 2024; Dorfner et al., 2024), our results are consistent with the observation that generalist models can perform competitively against domain-specialized models in medical tasks.

5.1.2. Influence of article access on performance

To assess the importance of context for critical appraisal, we evaluated models under three different settings as shown in [Figure 3](#). In addition to the question and answer choices, models were provided with either the full article, only the abstract, or no contextual information. The vast majority of questions in our dataset require access to the article to be answered correctly. While a few questions can be answered without reading the article, relying mostly on general medical knowledge, we still expect a performance drop when the article is not provided.

Performance varied significantly depending on the amount of context available. When given access to the full article, models achieved their highest scores; for example, GPT-4.1 reached an EMR of 0.49, and Qwen3-32B 0.36. While this confirms that access to the entire article improves question comprehension and accuracy, the overall performance remains moderate, reflecting the inherent difficulty of the task. Moreover, it is possible that some articles may have been included in the models' pre-training data. In such cases, the models could retrieve or reproduce information from memory rather than truly understanding the context, which might artificially inflate performance scores.

With only the abstract, model performance

Model	EMR	F1	Hamming	LCA
Random (1-3 options)	0.03	0.39	0.29	0.16
Only first-option	0.00	0.20	0.11	0.00
Most 2 frequent answers	0.03	0.45	0.33	0.18
Most 3 frequent answers	0.03	0.55	0.42	0.20
GPT4.1	0.49	0.84	0.78	0.68
GPT4o-mini	0.25	0.75	0.65	0.50
Qwen2.5-3B-Instruct	0.10	0.59	0.46	0.31
Qwen2.5-3B-GRPO-medical-reasoning	0.11	0.57	0.45	0.30
Qwen3-8B	0.19	0.68	0.57	0.42
ll-Medical-8B	0.13	0.51	0.43	0.30
Qwen3-32B	<u>0.37</u>	<u>0.78</u>	<u>0.70</u>	<u>0.58</u>
Gemma3-27B-text-IT	0.27	0.75	0.65	0.51
Medgemma-27B-text-IT	0.28	0.73	0.63	0.50

Table 4: Results of long-context models on the biomedical MCQA task of our dataset: zero-shot with full-text article provided. Bold values indicate the best results per metric, and underlined values the second-best. Baselines include random or frequency-based answer selection strategies.

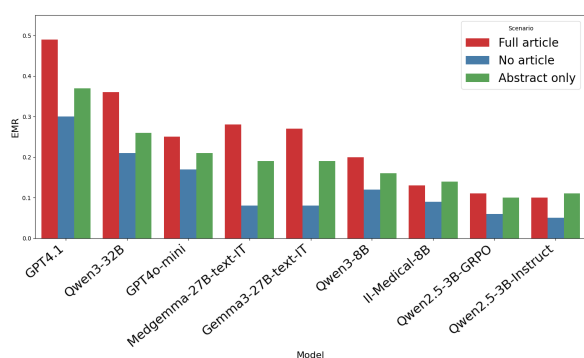


Figure 3: Exact Match Ratio comparison across different evaluation scenarios, illustrating model performance when provided with the full article, only the abstract, or no context (only the question and answer options with the instruction prompt).

dropped slightly but remained better than in the no-context condition. GPT-4o-mini showed slightly lower results with just the abstract compared to the full article, yet its performance remained competitive, suggesting that the abstract alone contains a substantial portion of the relevant information. Without any contextual input, all models experienced a notable decrease in performance, with EMR scores dropping by 5 to 15 points.

These findings highlight the importance of full-text access to maximize model performance on the task. While models like GPT-4.1 and Qwen3-32B demonstrate relative robustness even with limited or no context, by possibly mastering questions which do not require context, smaller models generally struggle to compensate for missing information.

We annotated a subset of 16 questions to indicate whether they require context to be answered correctly (field `requires_context`, can be `true` or `false`). These annotations were based on general trends observed in model performance:

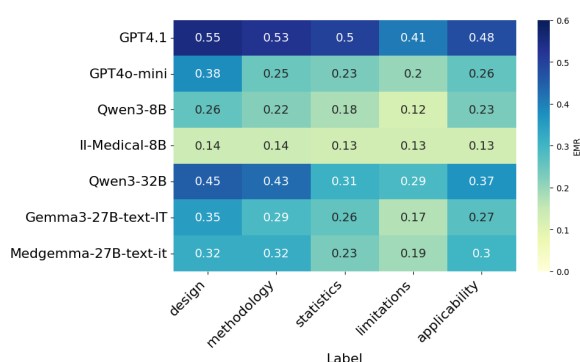


Figure 4: Heatmap of Exact Match Ratio by model and label for the MCQA task, illustrating performance differences across reasoning categories in the critical appraisal of scientific articles. Labels correspond to distinct cognitive skills required to answer the questions as described in Table 1.

whether models could answer questions correctly with or without access to the article.

Thanks to this subset, we can see some regularities and performance patterns with context or no context provided: some questions, particularly those directly related to specific aspects of the study (like *"This is a study of:"*) consistently require the article for all models. Other questions show little difference in performance between the with- and without-context settings, suggesting that context is not necessary for these. Conversely, some questions (like *"What was the main reason the data and safety monitoring board re-evaluated the sample size during the trial?"*) were answered more accurately without the article, suggesting the presence of a bias for this type of question rather than a benefit from using contextual information.

Model	EMR	F1	Hamming	LCA
Qwen3-8B (w/o→with)	0.19→0.35	0.68→0.75	0.57→0.66	0.42→0.55
Qwen3-32B (w/o→with)	0.37→0.45	0.78→0.81	0.70→0.73	0.58→0.64
GPT-OSS-20b (low→high)	0.36→0.49	0.77→0.81	0.60→0.75	0.57→0.66
GPT-OSS-120b (low→high)	0.46→0.54	0.81→0.85	0.74→0.79	0.65→0.71
GPT4.1 (w/o→with)	0.49→0.53	0.84→0.85	0.78→0.79	0.68→0.71

Table 5: Comparison of model performance according to "reasoning" tokens generation level on the CareMedEval benchmark.

5.1.3. Evaluation details by labels

Each question in our dataset was annotated with one or more labels highlighting the cognitive and analytical skills required to answer it. These labels allow us to analyze which types of questions are more or less challenging for the models. These results are presented in Figure 4.

Models appear to struggle the most with questions labeled *limitations*, which involve reviewing the biases or limitations of the study. This label typically requires fine-grained contextual understanding and often goes beyond what is explicitly stated in the text. The low scores suggest that models are subject to difficulty with implicit critical reasoning.

The *statistics* label, which requires understanding and interpreting statistical results, also shows lower performance compared to other categories. This can be partly explained by limitations in quantitative reasoning, but also by the fact that articles are provided in plain text format, excluding figures where statistical information is often presented.

Questions labeled *design* and *methodology* are those on which the models perform best. The top-performing models reach strong scores in these categories which may reflect the models' ability to recognize study structure and generalize research concepts as typically presented in medical articles.

5.1.4. Impact of reasoning tokens generation on performance

We evaluated the impact of explicit reasoning tokens generation on performance by comparing standard predictions (*without reasoning*) to tests where models were prompted or trained to produce a reasoning sequence before answering (*with reasoning*). For models that do not expose a mode without reasoning, we compare low and high reasoning effort presets. The results are available in Table 5. For each experiment we generated a single reasoning trace based on default parameters; on average, 879 tokens were generated ranging from 36 to 20,019 tokens across all models.

For GPT-4.1, reasoning is explicitly requested in the prompt, whereas for the Qwen3 models, we extract the reasoning part naturally generated by the model (content between `<think>` tags). We

also evaluate GPT-OSS models (OpenAI, 2025), according to low and high reasoning profiles. Even though we do not show results here, the medium reasoning profile yields results between those of the low and high profiles, generally closer to the later.

Incorporating intermediate reasoning steps improves performance over no/low reasoning across all metrics, sometimes substantially, suggesting that generating explicit reasoning tokens helps models produce more accurate answers. This result suggests that CareMedEval indeed requires some form reasoning and can evaluate efforts from LLM makers to address this aspect. We leave the manual evaluation of reasoning quality to future work as it is non trivial to match the reasoning trace output by models with the human-written justifications available in the dataset.

6. Conclusion

In this work, we introduced an original dataset designed for evaluation of critical appraisal of scientific articles in the medical domain, combining data collection and expert annotation. We assess the suitability of the dataset by computing the performance of a range of models.

Overall, our experiments show that larger models like GPT-4.1 and Qwen3-32B tend to perform better on the task and domain-specialized biomedical models do not reliably outperform generalist models, often showing similar levels of performance. However, using our LCA score based on real-life medical exams, we find that none of the tested models without reasoning achieve the passing score that human candidates typically reach.

Providing the full-text article as context considerably improves model performance compared to using only abstracts or no context at all. This underlines the necessity of access to complete scientific information for accurate question answering and critical reasoning. Moreover, allowing models to generate reasoning tokens improves performance, highlighting that reasoning is essential for producing more reliable and contextually grounded answers in a critical appraisal task.

In future work, we plan on extending the bench-

mark to vision LLMs that can leverage the content of figures which is sometimes referenced in questions or necessary to produce correct answers. We would also like to create an evaluation framework of the reasoning traces produced by models, compared to the justifications provided by experts.

7. Limitations & Ethics statement

Limitations Our study has several limitations that should be acknowledged.

First, our evaluation only used the textual content of the articles, without including figures, tables, or other materials found in the original PDFs. These parts often have important information that helps understand a study better, like how it was designed or its results.

Second, we evaluated a limited number of models, with a focus on general-purpose LLMs rather than domain-specific ones. Their performances may not fully reflect the potential of specialized biomedical models. In addition, our experiments used a fixed prompt structure without exploring prompt engineering variations or model-specific adaptations. Exploring alternative prompting strategies, including dynamic or adaptive prompts, could lead to improved performance.

Third, the manual annotation of justifications was performed by a small number of annotators, which may introduce variability and bias in the dataset. Expanding the number and diversity of annotators would help increase the reliability of the ground-truth justifications.

Fourth, the scientific articles included in our dataset are available on the internet and may have been part of the training data of the evaluated LLMs. This raises the possibility that some answers could benefit from memorization rather than genuine reasoning, potentially inflating performance metrics.

Finally, the dataset itself is relatively small (534 questions) and focused on a very specific educational context (French medical LCA exams). Future work could benefit from expanding the dataset in size and scope to cover a broader range of biomedical topics and question formats.

Ethics statement The goal of this work is to evaluate natural language processing technologies in order to better understand their capabilities and limitations, in particular within the context of the EU AI Act regulatory framework. We stress the importance of building useful tools for humans and the human society, that do not decrease expertise or agency, both of which are critical in the medical domain.

8. Acknowledgements

This work was financially supported by ANR MALADES (ANR-23-IAS1-0005). Experiments were conducted using HPC resources provided by the Laboratoire d'Informatique et Systèmes (LIS) in Marseille. We thank Guillaume Lomet, Elie Antoine and Irina Illina for helpful feedback on an earlier version of this work.

9. Bibliographical References

- Michael Alfertshofer, Cosima C. Hoch, Paul F. Funk, Katharina Hollmann, Barbara Wollenberg, Samuel Knoedler, and Leonard Knoedler. 2023. [Sailing the seven seas: A multinational comparison of ChatGPT's performance on medical licensing examinations](#). *Annals of Biomedical Engineering*, 52(6):1542–1545.
- Adrien Bazoge, Emmanuel Morin, Beatrice Daille, and Pierre-Antoine Gourraud. 2024. [Adaptation of biomedical and clinical pretrained models to french long documents: A comparative study](#).
- C. Glenn Begley and Lee M. Ellis. 2012. [Raise standards for preclinical cancer research](#). *Nature*, 483(7391):531–533.
- Iain Chalmers and Paul Glasziou. 2009. [Avoidable waste in the production and reporting of research evidence](#). *The Lancet*, 374(9683):86–89.
- Kay Dickersin. 1987. [Publication bias and clinical trials](#). *Controlled Clinical Trials*, 8(4):343–353.
- Felix J. Dorfner, Amin Dada, Felix Busch, Marcus R. Makowski, Tianyu Han, Daniel Truhn, Jens Kleesiek, Madhumita Sushil, Jacqueline Lammert, Lisa C. Adams, and Keno K. Bresslem. 2024. [Biomedical large languages models seem not to be superior to generalist models on unseen medical data](#).
- Jean-Baptist du Prel, Bernd Röhrig, and Maria Blettner. 2009. [Critical appraisal of scientific articles: part 1 of a series on evaluation of scientific publications](#). *Deutsches Ärzteblatt International*, 106(7):100–105.
- Google. 2025. [Medgemma-27b-text-it: Instruction-tuned medical llm based on gemma 3 \(27b\)](#). <https://huggingface.co/google/medgemma-27b-text-it>.
- Natalia Grabar, Vincent Claveau, and Clément Dalloux. 2018. [CAS: French Corpus with Clinical Cases](#). In *Ninth International Workshop on Health Text Mining and Information Analysis*

- (LOUHI) *Proceedings of the Workshop*, Ninth International Workshop on Health Text Mining and Information Analysis (LOUHI) Proceedings of the Workshop, pages 1–7, Bruxelles, France.
- Binglan Han, Teo Susnjak, and Anuradha Mathrani. 2024. [Automating systematic literature reviews with retrieval-augmented generation: A comprehensive overview](#). *Applied Sciences*, 14(19):9103.
- Jiawei He, Boya Zhang, Hossein Rouhizadeh, Yingjian Chen, Rui Yang, Jin Lu, Xudong Chen, Nan Liu, Irene Li, and Douglas Teodoro. 2025. [Retrieval-augmented generation in biomedicine: A survey of technologies, datasets, and clinical applications](#).
- Intelligent Internet. 2025. li-medical-8b. <https://huggingface.co/Intelligent-Internet/II-Medical-8B>.
- John P. A. Ioannidis. 2005. [Why most published research findings are false](#). *PLoS Medicine*, 2(8):e124.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2020. [What disease does this patient have? a large-scale open domain question answering dataset from medical exams](#).
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W. Cohen, and Xinghua Lu. 2019. [PubMedqa: A dataset for biomedical research question answering](#).
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#).
- Yanis Labrak, Adrien Bazoge, Richard Dufour, Mickael Rouvier, Emmanuel Morin, Béatrice Daille, and Pierre-Antoine Gourraud. 2023. [Frenchmedmcqa: A french multiple-choice question answering dataset for medical domain](#).
- Yanis Labrak, Adrien Bazoge, Oumaima El Khattari, Mickael Rouvier, Pacome Constant dit Beauvils, Natalia Grabar, Beatrice Daille, Solen Quiniou, Emmanuel Morin, Pierre-Antoine Gourraud, and Richard Dufour. 2024. [Drbenchmark: A large language understanding evaluation benchmark for french biomedical domain](#).
- Dacheng Li, Rulin Shao, Anze Xie, Ying Sheng, Lianmin Zheng, Joseph E. Gonzalez, Ion Stoica, Xuezhe Ma, and Hao Zhang. 2023. [How long can context length of open-source llms truly promise?](#) In *Workshop on Instruction Tuning and Instruction Following at NeurIPS 2023*. OpenReview preprint.
- Mingchen Li, Halil Kilicoglu, Hua Xu, and Rui Zhang. 2024a. [Biomedrag: A retrieval augmented large language model for biomedicine](#).
- Mingchen Li, Zaifu Zhan, Han Yang, Yongkang Xiao, Jiatao Huang, and Rui Zhang. 2024b. [Benchmarking retrieval-augmented large language models in biomedical nlp: Application, robustness, and self-awareness](#).
- Siru Liu, Allison B McCoy, and Adam Wright. 2025. [Improving large language model applications in biomedicine with retrieval-augmented generation: a systematic review, meta-analysis, and clinical development guidelines](#). *Journal of the American Medical Informatics Association*, 32(4):605–615.
- Xiangbin Meng, Xiangyu Yan, Kuo Zhang, Da Liu, Xiaojuan Cui, Yaodong Yang, Muhan Zhang, Chunxia Cao, Jingjia Wang, Xuliang Wang, Jun Gao, Yuan-Geng-Shuo Wang, Jia-ming Ji, Zifeng Qiu, Muzi Li, Cheng Qian, Tianze Guo, Shuangquan Ma, Zeying Wang, Zexuan Guo, Youlan Lei, Chunli Shao, Wenyao Wang, Haojun Fan, and Yi-Da Tang. 2024. [The application of large language models in medicine: A scoping review](#). *iScience*, 27(5):109713.
- Elliot Nelson, Georgios Kollias, Payel Das, Subhajit Chaudhury, and Soham Dan. 2024. [Needle in the haystack for memory based large language models](#).
- Anna Nikiforovskaya, Nikolai Kapralov, Anna Vlasova, Oleg Shpynov, and Aleksei Shpilman. 2020. [Automatic generation of reviews of scientific papers](#). In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 314–319.
- Aurélie Névéol, Cyril Grouin, Jeremy Leixa, Sophie Rosset, and Pierre Zweigenbaum. 2014. [The QUAERO French medical corpus: A resource for medical entity recognition and normalization](#). In *Proc of BioTextMining Work*, pages 24–30.
- OpenAI. 2025. [gpt-oss-120b & gpt-oss-20b model card](#).
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher

Berner, Lenny Bogdonoff, Oleg Boiko, Madeleine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth

Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nicolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#).

Amandine Quercia, Jamil Zagher, Christian Lovis, and Christophe Gaudet-Blavignac. 2024. [Medfrenchmark, a small set for benchmarking generative llms in medical french](#). *Studies in Health Technology and Informatics*, 316:601–605.

Shruti Singh, Nandan Sarkar, and Arman Cohan. 2024. [Scidqa: A deep reading comprehension dataset over scientific papers](#).

Chengrui Wang, Qingqing Long, Meng Xiao, Xunxin Cai, Chengjun Wu, Zhen Meng, Xuezhi Wang, and Yuanchun Zhou. 2024a. [Biorag: A rag-llm framework for biological question reasoning](#).

Jinqiang Wang, Huansheng Ning, Yi Peng, Qikai Wei, Daniel Tesfai, Wenwei Mao, Tao Zhu, and Runhe Huang. 2024b. [A survey on large language models from general purpose to medical applications: Datasets, methodologies, and evaluations](#).

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chu-jie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren

Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#).

Weizhe Yuan, Pengfei Liu, and Graham Neubig. 2022. [Can we automate scientific reviewing?](#) *Journal of Artificial Intelligence Research*, 73:12863–12894.

Hye Sun Yun, Iain J. Marshall, Thomas A. Trikalinos, and Byron C. Wallace. 2023. [Appraising the potential uses and harms of llms for medical systematic reviews](#).

Mingqian Zheng, Jiaxin Pei, Lajanugen Logeswaran, Moontae Lee, and David Jurgens. 2024. [When "a helpful assistant" is not really helpful: Personas in system prompts do not improve performances of large language models](#).