nature medicine

Consensus Statement

https://doi.org/10.1038/s41591-024-03425-5

The TRIPOD-LLM reporting guideline for studies using large language models

Received: 24 July 2024

Accepted: 21 November 2024

Published online: 8 January 2025

Check for updates

Jack Gallifant[1,2,3], Majid Afshar[4,29], Saleem Ameen[1,5,6,29], Yindalon Aphinyanaphongs[7,29], Shan Chen[3,8,29], Giovanni Cacciamani[9,10,29], Dina Demner-Fushman[11,29], Dmitriy Dligach[12,29], Roxana Daneshjou[13,14,29], Chrystinne Fernandes[1,29], Lasse Hyldig Hansen[1,15,29], Adam Landman[16,29], Lisa Lehmann[16,29], Liam G. McCoy[17,29], Timothy Miller[18,29], Amy Moreno[19,29], Nikolaj Munch[1,15,29], David Restrepo[1,20,29], Guergana Savova[18,29], Renato Umeton[21,29], Judy Wawira Gichoya[22,29], Gary S. Collins[23,24], Karel G. M. Moons[25,26], Leo A. Celi[1,27,28] & Danielle S. Bitterman[3,8] ✉

Large language models (LLMs) are rapidly being adopted in healthcare, necessitating standardized reporting guidelines. We present transparent reporting of a multivariable model for individual prognosis or diagnosis (TRIPOD)-LLM, an extension of the TRIPOD + artificial intelligence statement, addressing the unique challenges of LLMs in biomedical applications. TRIPOD-LLM provides a comprehensive checklist of 19 main items and 50 subitems, covering key aspects from title to discussion. The guidelines introduce a modular format accommodating various LLM research designs and tasks, with 14 main items and 32 subitems applicable across all categories. Developed through an expedited Delphi process and expert consensus, TRIPOD-LLM emphasizes transparency, human oversight and task-specific performance reporting. We also introduce an interactive website (https://tripod-llm.vercel.app/) facilitating easy guideline completion and PDF generation for submission. As a living document, TRIPOD-LLM will evolve with the field, aiming to enhance the quality, reproducibility and clinical applicability of LLM research in healthcare through comprehensive reporting.

Healthcare's embrace of large language models (LLMs) shows no signs of slowing down, with current and future deployment being considered in several domains across administrative and healthcare delivery use cases, including generating drafts for patient communication, medical document summarization, question answering, information retrieval, medical diagnosis, treatment recommendations, patient education and medical education[1–5]. The rapid advancements in LLMs have stretched existing regulatory and governance structures to their limits, exposing a patchwork of solutions that do not fully encompass the nuances of these all-purpose models[6–8]. More broadly, the speed of development of LLMs poses a challenge to journal and peer-review publication timelines and to regulatory agencies, who seek to provide timely guidance. To keep pace, researchers publish preprints quickly and take an ad hoc approach to reporting.

Reporting guidelines provide a scalable method for standardizing research, transparent reporting and the peer-review process. The TRIPOD (transparent reporting of a multivariable model for individual prognosis or diagnosis) initiative is a critical example that was first introduced in 2015 to establish minimum reporting standards for diagnostic and prognostic prediction model studies (https://www.tripod-statement.org)[9]. TRIPOD is one of the core guidelines for the EQUATOR (enhancing the quality and transparency of health research) network, which is an international effort that promotes transparent, accurate reporting of health research literature[10]. TRIPOD is widely endorsed and

A full list of affiliations appears at the end of the paper. ✉e-mail: dbitterman@bwh.harvard.edu

recommended by journals and is often included in journal instructions to authors. TRIPOD has subsequently been updated to incorporate best practices for artificial intelligence (AI) due to the substantially evolved machine learning landscape, resulting in TRIPOD + AI[11]. This is in addition to other guidelines that offer complementary guidance on AI development throughout the model life cycle[12–14].

LLMs represent a distinct frontier within AI, introducing unique challenges and considerations not fully addressed by original TRIPOD guidelines or their newer extensions as we shift from classifier AI models to generative AI. Here we report the TRIPOD-LLM statement, an extension of TRIPOD + AI[11], developed to address these unmet needs and designed to be a living checklist to nimbly adapt to the rapidly evolving field. This expansion broadens the TRIPOD scope beyond its original focus on prediction models and reflects the pervasive impact of LLMs across diverse areas of medical research and practice, from diagnosis to document summarization.

Rationale for TRIPOD-LLM

LLMs, as generative AI models, are autoregressive, meaning — in the simplest possible terms — that they are trained to predict the next word in a sequence given the words that preceded it. Yet this foundational training has been shown to equip them with capabilities to perform a wide range of healthcare-related natural language processing (NLP) tasks from a single model. This adaptability is commonly achieved through supervised fine-tuning or few-shot learning methods, which allow LLMs to handle new tasks with minimal examples[15,16]. Chatbot solutions (for example, ChatGPT) use LLMs as their foundation, upon which two more components are added: question answering (referred to as instruction tuning or supervised fine-tuning) and preference ranking (referred to as alignment). The unique methodological processes involved in LLMs and chatbots are not captured by current reporting guidelines, such as the choice of hyperparameters used for supervised fine-tuning, the intricacies of prompting, variability in model predictions, methods in evaluating natural language outputs and preference-based learning strategies — which require specific guidance and have a substantial impact on model reliability. In addition, the generalist and generative nature of LLMs requires more detailed guidance than covered in prior guidelines. Because LLMs can be applied to a broad range of use cases for which they were not specifically trained and were not necessarily represented in training data (for example, disease prevalence typically captured in a task-specific model's training data for a given use case), they require unique task-specific guidance for robust reporting and downstream reliability and safety.

The selection of appropriate automated and human metrics by which to evaluate generative output remains an open question, and, currently, a wide range of methodologies are applied to capture various facets of performance. For tasks where the output is truly unstructured text and cannot be resolved to a structured label, such as letter generation or summarization tasks, evaluation is particularly complex. In these cases, most automated metrics prioritize overlap and similarity between input and output text, producing scores that may not capture factual accuracy or relevance of the text produced, and may not detect hallucinations or omissions[17–19]. These scores reflect the degree of structural and lexicographical resemblance to reference texts, which, although important, capture only a fraction of what constitutes a comprehensive evaluation of performance and safety. Human evaluation of text is a subjective process, complicated by the ambiguity of language and uncertainties inherent to many clinical tasks. These challenges are heightened in medicine, where there is often no single correct answer and both aleatoric and epistemic uncertainty are common. Therefore, specific details are needed to guide reporting of how performance is evaluated. In this paper, we use the term LLM to refer to both LLMs and chatbots. Table 1 highlights key categories of tasks applicable to the healthcare domain and provides notable definitions and examples of existing relevant work.

The new complexities introduced by LLMs include concerns regarding hallucinations, omissions, reliability, explainability, reproducibility, privacy and biases being propagated downstream, which can adversely affect clinical decision-making and patient care[20–26]. Furthermore, growing partnerships between electronic health record (EHR) vendors, technology companies and healthcare providers have led to deployment horizons that far outpace current regulatory timelines[8,27]. To safeguard LLM use and increase transparency, standardization in the development and reporting of LLMs is essential — to ensure consistency, reliability and verifiability, akin to established clinical-grade evaluation in other scientific domains[28–30].

The TRIPOD-LLM statement

The TRIPOD-LLM comprises a checklist of items considered essential for good reporting of studies that are developing, tuning, prompt engineering or evaluating an LLM (Table 2). Box 1 summarizes noteworthy additions and changes to TRIPOD-2015 and TRIPOD + AI, with key definitions provided in Box 2. The TRIPOD-LLM checklist comprises 19 main items about the title (one item), abstract (one item), introduction (two items), methods (eight items), open science practices (one item), patient and public involvement (one item), results (three items) and discussion (two items). These main items are further divided into 50 subitems. Of these, 14 main items and 32 subitems are applicable to all research designs and LLM tasks. The remaining 5 main items and 18 subitems are specific to particular research designs or LLM task categories. As discussed in the Methods, the TRIPOD-LLM statement introduces a modular format given the varied nature of LLM studies (Table 1), where some items are only relevant for specific research designs and LLM task categories. These designs and task categories are broad but not mutually exclusive, depending on the context of a specific study, and may need to evolve as the application of LLMs evolves. A separate checklist for journal or conference abstracts of LLM-based studies is included, and the TRIPOD + AI for abstracts statement[18] is revised (TRIPOD-LLM for abstracts; Table 3), reflecting new content and maintaining consistency with TRIPOD-LLM.

The recommendations contained within TRIPOD-LLM are for completely and transparently reporting on how LLM-based research was conducted; TRIPOD-LLM does not prescribe how to develop or evaluate LLMs specifically. The checklist is not a quality appraisal tool. Similarly, CANGARU (ChatGPT, generative artificial intelligence and natural large language models for accountable reporting and use)[31] and CHART (Chatbot Assessment Reporting Tool)[32] are complementary guidelines that relate to generative AI more broadly and chatbots specifically.

In addition to the TRIPOD website (https://www.tripod-statement.org), an accompanying interactive website was developed (https://tripod-llm.vercel.app/) to present the required questions based on research design(s) and task(s) for ease of completion. This site can be used to render a final PDF suitable for submission. Fillable templates for the TRIPOD-LLM checklist can also be downloaded from https://www.tripod-statement.org. News, announcements and information relating to TRIPOD-LLM and the release of subsequent statements can be found on the TRIPOD-LLM website and TRIPOD website (https://www.tripod-statement.org).

An example of a completed TRIPOD-LLM checklist for a previously published study reporting the pretraining, fine-tuning, retrospective evaluation and clinical deployment of an LLM for clinical and operational hospital tasks is presented in Supplementary Table 1 (ref. 6). A fillable checklist is provided in Supplementary Table 2, and an explanation and elaboration document is available in Supplementary Table 3 to aid understanding of new items.

TRIPOD-LLM statement as a living document

Given the rapid pace of the field and the timeline for interaction with healthcare workers and patients, the decision was made to create an accelerated TRIPOD-LLM statement to provide timely guidance for LLM

Table 1 | Research design and LLM task categories for the modular TRIPOD-LLM guidelines

Task	Definition	Example
Research design		
De novo LLM development	Building a new language model from scratch or substantially fine-tuning existing base models to develop new functionalities or to adapt to new tasks.	A study pretraining a new LLM on a hospital's clinical data, for example, ref. 56
LLM methods	Quantitative or theoretical investigations that focus on new architectures of model design, new computational methods to understand LLMs, new methods to evaluate LLMs and/or new methods to optimize LLM prompting.	A study of a new retrieval-augmented generation LLM framework for medicine, for example, ref. 57
LLM evaluation	Assessing or testing an existing LLM to determine its efficacy, accuracy or suitability for a specific task within healthcare; may also include evaluating the risks and biases arising from using it.	A study investigating biased diagnostic reasoning in an existing LLM, for example, ref. 21
LLM evaluation in healthcare settings	Evaluating an LLM when used as part of a clinical workflow, focusing on its integration and impact on clinical, administrative or workforce outcomes.	A study reporting the performance of an LLM deployed in real-time to predict outcomes in hospitalized patients, for example, ref. 6
LLM task		
Text processing	Manipulating and lower-order processing of text data, which includes tasks including but not limited to tokenization, parsing and entity recognition.	A study investigating a new LLM approach to named entity recognition, for example, ref. 58
Classification	Assigning predefined labels to text data.	A study fine-tuning an LLM to determine whether or not a sentence in a clinic note mentions one or several social determinants of health, for example, ref. 59
Long-form question answering	Providing detailed answers to complex queries, which can involve reasoning over multiple documents or pieces of evidence. Please note that multichoice question answering is subsumed under classification.	A study investigating the ability of an existing LLM to respond to patient portal messages, for example, ref. 24
Information retrieval	The process of fetching relevant information from large datasets based on specific queries, which is relevant for tasks such as literature review or patient history retrieval.	A study that trained a transformer model to retrieve biomedical publications relevant to a user's query, for example, ref. 60
Conversational agent (chatbot)	Responding to and engaging in conversation with users — often used for patient interaction, health advisories or as virtual assistants for healthcare providers.	A study investigating whether access to an interactive LLM-based chatbot impacts clinicians' diagnostic reasoning, for example, ref. 61
Documentation generation	Automatically creating medical documentation from clinical data, dictations or recordings.	A study evaluating the quality of clinical notes automatically generated from ambient clinic recordings, for example, ref. 5
Summarization and simplification	Condensing large text documents into shorter versions or simplifying the content for easier comprehension is useful in patient education or in creating executive summaries of medical records.	A study evaluating the ability of LLMs to convert discharge summaries into patient-friendly plain language, for example, ref. 62
Machine translation	Converting text from one language to another.	A study comparing the ability of smaller language models fine-tuned for translation versus generalist LLMs to translate biomedical text between Spanish and English, for example, ref. 63
Outcome forecasting	Predicting future medical outcomes based on historical data, which can be used in prognosis estimation or treatment effectiveness studies.	A study investigating the ability of LLMs to predict out-of-hospital mortality in patients admitted to intensive care units, for example, ref. 64

use in (bio)medical and other healthcare applications. This guidance has been designed as part of a living document hosted on an interactive website to facilitate agile versioning, refinement from user testing, updates as the field evolves and regular meetings to intake and evaluate new standards. Thus, as the reporting recommendations are anticipated to evolve, users are directed to the most current version of the guidelines at https://tripod-llm.vercel.app/.

Our approach to the living TRIPOD-LLM statement is informed by processes established in developing living systematic reviews[33,34] and clinical practice guidelines[35,36], which have been adopted to address a similar need to provide updated, timely recommendations based on evolving evidence. Public comments on the statement will be collected from the community via multiple avenues to enhance accessibility — a project-specific GitHub repository, the TRIPOD-LLM website and the main TRIPOD website (https://www.tripod-statement.org/). We encourage input both on usability, such as language that may be ambiguous or redundant, and on the content of the guidelines themselves. As a few examples, users may suggest a change to an item to make it more feasible in practice, recommend a new item be added, recommend adding or removing items assigned to a given research design of the LLM task module or recommend changes to the research design or LLM task module categories.

An expert panel will convene every 3 months to discuss updates. Before the meeting, members will review the intercurrent literature to inform any updates. The units for update will be checklist items, research design categories and LLM task categories delineated in the statement. At the meeting, the panel will discuss the current statement and suggest revisions considering public comments, literature review and subject matter expertise. The steering committee will revise the statement based on this discussion, and this will be circulated to the expert panel for final review and approval. Review can result in the following action for each component of the TRIPOD-LLM statement (adapted from ref. 33) — items, research designs and LLM tasks:

1. No modification.
2. Modification of substantive content (small, editorial revisions such as rewording for clarity and and correcting types will not be considered a modification).
3. Merging of one or more components together (merging will only take place within a component type).
4. Splitting one component into two or more components (splitting will only take place within a component type).
5. Retiring the component from the statement.

https://doi.org/10.1038/s41591-024-03425-5

Table 2 | TRIPOD-LLM checklist

Section	Item	Description	Research design	LLM task
Title	1	Identify the study as developing, fine-tuning and/or evaluating the performance of an LLM, specifying the task, the target population and the outcome to be predicted.	All	All
Abstract	2	See TRIPOD-LLM for abstracts.	All	All
Introduction				
Background	3a	Explain the healthcare context/use case (for example, administrative, diagnostic, therapeutic and clinical workflow) and rationale for developing or evaluating the LLM, including references to existing approaches and models.	All	All
	3b	Describe the target population and the intended use of the LLM in the context of the care pathway, including its intended users in current gold standard practices (for example, healthcare professionals, patients, public or administrators).	E H	All
Objectives	4	Specify the study objectives, including whether the study describes the initial development, fine-tuning or validation of an LLM (or multiple stages).	All	All
Methods				
Data	5a	Describe the sources of data separately for the training, tuning and/or evaluation datasets and the rationale for using these data (for example, web corpora, clinical research/trial data, EHR data or unknown).	All	All
	5b	Describe the relevant data points and provide a quantitative and qualitative description of their distribution and other relevant descriptors of the dataset (for example, source, languages and countries of origin).	All	All
	5c	Specifically state the date of the oldest and newest item of text used in the development process (training, fine-tuning and reward modeling) and the evaluation datasets.	All	All
	5d	Describe any data preprocessing and quality checking, including whether this was similar across text corpora, institutions and relevant sociodemographic groups.	All	All
	5e	Describe how missing and imbalanced data were handled and provide reasons for omitting any data.	All	All
Analytical methods	6a	Report the LLM name, version and last date of training.	All	All
	6b	Report details of the LLM development process, such as LLM architecture, training, fine-tuning procedures and alignment strategy (for example, reinforcement learning and direct preference optimization) and alignment goals (for example, helpfulness, honesty and harmlessness).	M D	All
	6c	Report details of how the text was generated using the LLM, including any prompt engineering (including consistency of outputs), and inference settings (for example, seed, temperature, max token length and penalties), as relevant.	M D E	All
	6d	Specify the initial and postprocessed output of the LLM (for example, probabilities, classification and unstructured text).	All	All
	6e	Provide details and rationale for any classification and, if applicable, how the probabilities were determined and thresholds identified.	All	C OF
LLM output	7a	Include metrics that capture the quality of generative outputs, such as consistency, relevance, accuracy and presence/type of errors compared to gold standards.	All	QA IR DG SS MT
	7b	Report the outcome metrics' relevance to the downstream task at deployment time and, where applicable, the correlation of metric to human evaluation of the text for the intended use.	E H	All
	7c	Clearly define the outcome, how the LLM predictions were calculated (for example, formula, code, object and API), the date of inference for closed-source LLMs and evaluation metrics.	E H	All
	7d	If outcome assessment requires subjective interpretation, describe the qualifications of the assessors, any instructions provided, relevant information on demographics of the assessors and interassessor agreement.	All	All
	7e	Specify how performance was compared to other LLMs, humans and other benchmarks or standards.	All	All
Annotation	8a	If annotation was done, report how the text was labeled, including providing specific annotation guidelines with examples.	All	All
	8b	If annotation was done, report how many annotators labeled the dataset(s), including the proportion of data in each dataset that was annotated by more than one annotator, and the interannotator agreement.	All	All
	8c	If annotation was done, provide information on the background and experience of the annotators or the characteristics of any models involved in labeling.	All	All
Prompting	9a	If research involved prompting LLMs, provide details on the processes used during prompt design, curation and selection.	All	All
	9b	If research involved prompting LLMs, report what data were used to develop the prompts.	All	All

Table 2 (continued) | TRIPOD-LLM checklist

Section	Item	Description	Research design	LLM task
Summarization	10	Describe any preprocessing of the data before summarization.	All	SS
Instruction tuning/alignment	11	If instruction tuning/alignment strategies were used, what were the instructions, data and interface used for evaluation, and what were the characteristics of the populations doing the evaluation?	M D	All
Compute	12	Report compute, or proxies thereof (for example, time on what and how many machines, cost on what and how many machines, inference time, floating-point operations per second), required to carry out methods.	M D E	All
Ethical approval	13	Name the institutional research board or ethics committee that approved the study and describe the participant-informed consent or the ethics committee waiver of informed consent.	All	All
Open science	14a	Give the source of funding and the role of the funders for the present study.	All	All
	14b	Declare any conflicts of interest and financial disclosures for all authors.	All	All
	14c	Indicate where the study protocol can be accessed or state that a protocol was not prepared.	H	All
	14d	Provide registration information for the study, including register name and registration number, or state that the study was not registered.	H	All
	14e	Provide details of the availability of the study data.	All	All
	14f	Provide details of the availability of the code to reproduce the study results.	All	All
Public involvement	15	Provide details of any patient and public involvement during the design, conduct, reporting, interpretation or dissemination of the study or state no involvement.	H	All
Results				
Participants	16a	When using patient/EHR data, describe the flow of text/EHR/patient data through the study, including the number of documents/questions/participants with and without the outcome/label and follow-up time as applicable.	E H	All
	16b	When using patient/EHR data, report the characteristics overall and for each data source or setting and development/evaluation splits, including the key dates, key characteristics and sample size.	E H	All
	16c	For LLM evaluation that includes clinical outcomes, show a comparison of the distribution of important clinical variables that may be associated with the outcome between development and evaluation data, if available.	E H	All
	16d	When using patient/EHR data, specify the number of participants and outcome events in each analysis (for example, for LLM development, hyperparameter tuning and LLM evaluation).	E H	All
Performance	17	Report LLM performance according to prespecified metrics (see item 7a) and/or human evaluation (see item 7d).	All	All
LLM updating	18	If applicable, report the results from any LLM updating, including the updated LLM and subsequent performance.	All	All
Discussion				
Interpretation	19a	Give an overall interpretation of the main results, including issues of fairness in the context of the objectives and previous studies.	All	All
Limitations	19b	Discuss any limitations of the study and their effects on any biases, statistical uncertainty and generalizability.	All	All
Usability of the LLM in context	19c	Describe any known challenges in using data for the specified task and domain context with reference to representation, missingness, harmonization and bias.	E H	All
	19d	Define the intended use for the implementation under evaluation, including the intended input, end-user and level of autonomy/human oversight.	E H	All
	19e	If applicable, describe how poor quality or unavailable input data should be assessed and handled when implementing the LLM; that is, what is the usability of the LLM in the context of current clinical care.	E H	All
	19f	If applicable, specify whether users will be required to interact in the handling of the input data or use of the LLM, and what level of expertise is required of users.	E H	All
	19g	Discuss any next steps for future research, with a specific view of the applicability and generalizability of the LLM.	All	All

For studies using existing LLMs, users should include reference(s) to reportable information if provided by the original developers or state that this information is not available. M, LLM methods; D, de novo LLM development; E, LLM evaluation; H, LLM evaluation in healthcare settings; C, classification; OF, outcome forecasting; QA, long-form question answering; IR, information retrieval; DG, document generation; SS, summarization and simplification; MT, machine translation; API, application programming interface.

Release of a new version of the statement will be disseminated to the community through postings on the TRIPOD-LLM website, the main TRIPOD website (https://www.tripod-statement.org/), the EQUATOR Network website (https://www.equator-network.org/reporting-guidelines/) and postings on social media accounts. Emails will be sent to journal editors to inform them of the update and ensure that author instructions refer to the most current versioning. Users are requested to cite the version of the statement used in their research.

At each review meeting, the membership of the expert panel will be reviewed for expertise, diversity and representation, and new members

BOX 1

TRIPOD-LLM noteworthy changes and additions to TRIPOD-2015 and TRIPOD + AI

New checklist for reporting on LLMs. A dedicated checklist has been developed to address the unique aspects of reporting LLMs, reflecting their distinct characteristics and the specific methodologies they use compared to other AI and prediction models.

Living guideline. The checklist is designed as a living document, which will be updated on a regular basis based on a review of the literature and input from the community. This approach was taken due to the rapid developments in the field, enabling agile versioning, refinement from user testing and timely updates as the field advances.

Task-specific guidance. The checklist includes a new section that provides task-specific guidance designed to address the particular challenges and needs associated with different LLM applications in healthcare. This addition ensures that reporting is tailored and relevant to the specific functions and objectives of the LLM under study.

Enhanced emphasis on transparency and fairness. The new guidelines emphasize 'transparency' and 'fairness', highlighting the importance of recognizing and addressing societal biases that may be encoded in clinical models. The checklist integrates these concepts throughout, ensuring that bias and fairness are considered at every stage of the model's life cycle.

Modular framework. The new guidelines are modular, with different requirements based on the research design(s) and LLM task(s) that are reported in a given study. This change was motivated by the wide variety of applications and approaches in biomedical LLM research, from model development through evaluation, necessitating more specialized reporting items.

will be solicited if and when gaps are identified. Expert panel members will also have the authority to trigger an ad hoc review of the guidelines to accommodate major, unexpected changes in the field that warrant more urgent discussion.

Discussion

TRIPOD-LLM has been developed to guide researchers, journals, healthcare professionals, LLM developers (commercially and non-commercially) and healthcare institutions in the rapidly evolving field of biomedical and healthcare LLMs. It represents minimum reporting recommendations for studies describing LLMs' development, tuning or evaluation. Reporting TRIPOD-LLM items will help users understand and appraise the quality of LLM study methods, increase transparency around study findings, reduce overinterpretation of study findings, facilitate replication and reproducibility and aid in implementing the LLM.

Transparency throughout the model life cycle has been strongly emphasized in the guidelines. Detailed documentation is emphasized at each stage of an LLM's life cycle[37]; for example, during the development and fine-tuning phases, there is an emphasis on disclosing the origins and processing of training data. Moreover, the LLM version and specifics of any fine-tuning or alignment modeling processes on top of existing foundation models must be transparently reported to enable fair comparisons of LLMs. This includes specifying the cut-off dates for when training data were collected to clarify the temporal relevance of training datasets and the potential for data leakage or contamination during evaluation. In addition, studies should document the model version date and whether the model was frozen or remained dynamic during the data collection phase. Transparency regarding input data is essential because LLMs are typically trained on multiple public large-scale datasets — and thus inherently risk incorporating existing societal biases and inequities in the form of stigmatizing language, as well as statistical risk allocation in disparate groups, necessitating a comprehensive and transparent approach to curating data sources and understanding potential biases[21,22,38–41].

Human insight and oversight are critical components of the TRIPOD-LLM statement, reflecting an emphasis on components eventually critical for the responsible deployment of LLMs (although deployment reliability and observability are outside the scope of this paper)[42–44]. The guidelines include requirements for increased reporting of the expected deployment context and specifying the levels of autonomy assigned to the LLM, if applicable. Furthermore, there is a focus on the quality control processes used in dataset development and evaluation, such as qualifications of human assessors, requirements for dual annotation and specific details on instructions provided to assessors to ensure that nuances of text evaluation are captured, thus facilitating reliable assessments of safety and performance.

Prompting and task-specific performance are key additions necessitated by the unique characteristics of LLMs. The variability in prompt engineering approaches can have a considerable influence on LLM performance, potentially skewing benchmark comparisons and real-world applicability[45,46]. Where relevant, reports must include comprehensive descriptions of data sources used for developing prompts, LLM model names and versions, any preprocessing steps undertaken and methods used in prompt engineering. This ensures that prompts are effectively designed to elicit stable and reproducible performance from LLMs. Additionally, the guidelines call for clear reporting on evaluation settings, including instructions and interfaces used and characteristics of populations involved in evaluations. This is intended to emphasize whether LLM performance is assessed under conditions that closely mimic real-world applications, providing a reliable measure of its practical utility.

We anticipate that key users and beneficiaries of TRIPOD-LLM will be (1) academic and industry researchers authoring papers, (2) journal editors and peer reviewers evaluating research papers, and (3) other stakeholders (for example, the research community in general, academic institutions, policymakers, funders, regulators, patients, study participants, industry and the broader public) who will benefit from increased transparency and quality of LLM research. We encourage editors, publishers and the industry more broadly to support adherence to TRIPOD-LLM by referring to a link within the journal's instructions to authors, enforcing its use during the submission and peer-review process and making adherence to the recommendations an expectation. We also encourage funders to require applications for LLM studies to include a plan to report their model according to the TRIPOD-LLM recommendations, thereby minimizing research waste and ensuring value for money.

Of note, this guideline was developed with text-only LLMs in mind; however, advances in multimodal models incorporating LLMs, such as vision-language models[47], are now rapidly emerging — illustrating the need for rapid, nimble approaches for reporting guidelines. Many of the reporting considerations will be shared between text-only LLMs and these multimodal models. For example, for vision-language models, both text and image preprocessing should be reported. However, unique considerations may arise that merit discussion in future versioning of TRIPOD-LLM or related guidelines. For example, studies

BOX 2

Glossary of terms

Please note the definitions and descriptions given relate to the specific context of TRIPOD-LLM and the use of the terms in the guideline. They are not necessarily applicable to other areas of research.

Attention mechanism: A component in neural networks that allows the model to focus on different parts of the input when producing each part of the output, crucial for handling long-range dependencies in sequence data.

Chain-of-thought prompting: A prompting technique that encourages the model to break down complex reasoning tasks into step-by-step thought processes, often improving performance on logical and mathematical tasks.

Confabulation: An alternative term for hallucination, highlighting the production of false information that is not intended to deceive.

Data leakage: The use of test data during model training and/or fine-tuning, resulting in a model that performs better than the corresponding performance on unseen data.

Decoder: A component of a model that converts vectorized input data back into a text sequence.

Autoregressive model: A type of transformer-based model that predicts the next component in a sequence, for example, the next word in a sentence, based on the preceding sequence. Current state-of-the-art LLMs, including generative pretrained transformers, are autoregressive models.

Embedding: A representation of text in a high-dimensional vector space where similar words have similar representation, capturing semantic meaning (see also Vector).

Encoder: A component of a model that processes the input data, transforming it into a vectorized format or representation that the model can understand.

Encoder–decoder: A model architecture framework combining an encoder and decoder to transform input data into an output.

Few-shot learning: The model learns to perform a task proficiently with a very small number of exemplars. In some cases, the number of exemplars is specified in place of 'few', such as one-shot learning.

Fine-tuning: A process where a pretrained model is further trained on a smaller, domain-specific dataset to specialize its knowledge for specific tasks.

GPT: A family of autoregressive transformer-based models for natural language understanding and generation. These models are pretrained to predict the next word in a sentence.

Hallucination: A phenomenon where a language model generates text that is unrelated or loosely related to the input data, often manifesting as fabrications or inaccuracies.

In-context learning: The ability of a model to learn a new task from examples provided within the prompt at inference time.

Instruction tuning: A fine-tuning approach where models are trained on datasets consisting of natural language instructions and their corresponding desired outputs, improving the model's ability to follow diverse instructions.

Prompt: The query or instruction that is input into an LLM to elicit a response.

Reinforcement learning: A machine learning technique commonly used in LLM development that trains a model to optimize its output according to humans' preferences by providing rewards in response to actions.

Prompt engineering: A process that guides models to generate desired outputs. Examples include prompt iterations, prompting with examples, and chain-of-thought prompting.

Retrieval-augmented generation: A technique that combines information retrieval from an external knowledge base with text generation, allowing LLMs to access and incorporate up-to-date or domain-specific information.

Temperature: A parameter that controls the randomness of predictions by scaling the logits before applying softmax, affecting the diversity of generated text.

Tokenization: The process of converting text into smaller units, such as words or phrases, to facilitate their processing in NLP tasks.

Transformer: A commonly used neural network architecture that has advanced the field of NLP. Unlike its predecessors, the transformer processes sequences of data in parallel using self-attention mechanisms, improving efficiency and the ability to capture complex dependencies within the text.

Vector: A numerical representation of data. In LLMs, text data is represented as a type of vector known as contextual embeddings, where the vector for each token (word piece) is influenced by the surrounding words.

Zero-shot learning: The ability of a model to correctly perform tasks it has never explicitly been trained to do, based on its understanding and generalization capabilities.

reporting the development of LLMs that use imaging data should report details of image acquisition. In the interim, we suggest that studies reporting the development and/or evaluation of a method that includes an LLM as a primary component use the TRIPOD-LLM statement, although we acknowledge this may be subject to interpretation. We advise that users keep in mind the goals of reproducibility, understandability and transparency to take a common-sense approach to deciding on the appropriate reporting guideline and to interpreting

Table 3 | TRIPOD-LLM for abstracts

Section	Item	Checklist item	Research design	LLM task
Title	2a	Identify the study as developing, fine-tuning and/or evaluating the performance of an LLM, specifying the task, the target population and the outcome to be predicted.	All	All
Background	2b	Provide a brief explanation of the healthcare context, use case and rationale for developing or evaluating the performance of an LLM.	E H	All
Objectives	2c	Specify the study objectives, including whether the study describes LLMs development, tuning and/or evaluation	All	All
Methods	2d	Describe the key elements of the study setting.	All	All
	2e	Detail all data used in the study, specify data splits and any selective use of data.	M D E	All
	2f	Specify the name and version of LLM(s) used.	All	All
	2g	Briefly summarize the LLM-building steps, including any fine-tuning, reward modeling and RLHF.	M D	All
	2h	Describe the specific tasks performed by the LLMs (for example, medical QA, summarization and extraction), highlighting key inputs and outputs used in the final LLM.	All	All
	2i	Specify the evaluation datasets/populations used, including the endpoint evaluated, and detail whether this information was held out during training/tuning where relevant and what measure(s) were used to evaluate LLM performance.	All	All
Results	2j	Give an overall report and interpretation of the main results.	All	All
Discussion	2k	Explicitly state any broader implications or concerns that have arisen in light of these results.	All	All
Other	2l	Give the registration number and name of the registry or repository (if relevant).	H	All

RLHF, reinforcement learning with human feedback.

the relevant components of TRIPOD-LLM statement items to report multimodal LLMs. Users may also refer to methodological guides from multiple AI fields, such as radiomics[48,49], to inform their reporting.

Assurance labs, such as the Coalition for Health AI[50] and Epic AI Labs[51,52], or internal validation standards are expected to be of importance in the generation, verification, certification and maintenance of model cards ('at a glance' summary of model metadata) for clinical AI. It is our opinion that the TRIPOD-LLM standard can and should inform approaches by these labs to assure LLMs in ways that meet the required regulatory bar for AI (for example, the Biden administration's 'Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence', the US AI Safety Institute[53], the US Office of the National Coordinator for Health Information Technology HTI-1 Final Rule[54] and the European Union AI Act[55]) and build confidence in patients, clinician and other stakeholders about the utility and trustworthiness of clinical AI.

It must also be emphasized that LLM evaluation and validation require specialized expertise and resources. To ensure equitable and safe deployment, investments into LLM development should be balanced by investments into infrastructure that enables robust validation beyond large academic settings. Moreover, this checklist should be seen as part of a continuous process for evaluating LLMs due to the temporal and geographic-specific contexts these models inherit, which can impact the generalizability of performance and fairness across sites or at the same site over time. These shifts can be even more unpredictable than traditional ML models due to their user-dependent nature, and thus, considerable effort must be placed on understanding trends and heterogeneity of effects instead of single-point estimates that proclaim universal validation.

Limitations of the current TRIPOD-LLM checklist arise from the nearly unprecedented speed of development and publishing in this space, necessitating rapid guideline development to provide guidance to the research community. We carried out an expedited Delphi process to arrive at the initial version of the checklist included here but acknowledge that this may introduce limitations in consensus and breadth of input. We have therefore implemented the living statement approach to enable more nimble incorporation of feedback and adaptation to rapidly evolving methods. The living statement necessarily means that the checklist included in this manuscript will likely become outdated; users are directed to https://tripod-llm.vercel.app/, which will be kept up to date with the most current version.

Conclusion

TRIPOD-LLM aims to assist authors in the complete reporting of their study and help LLM developers, researchers, peer reviewers, editors, policymakers, end-users (for example, healthcare professionals) and patients understand data, methods, findings and conclusions of LLM-driven research. Adhering to the TRIPOD-LLM reporting recommendations may promote the best and most efficient use of research time, effort and money, enhancing the value of LLM research to maximize positive impact.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer-review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41591-024-03425-5.

References

1. Chen, Z. et al. MEDITRON-70B: scaling medical pretraining for large language models. Preprint at https://doi.org/10.48550/arXiv.2311.16079 (2023).
2. OpenAI, Josh, A. et al. GPT-4 technical report. Preprint at https://doi.org/10.48550/arXiv.2303.08774 (2023).
3. Singhal, K. et al. Large language models encode clinical knowledge. *Nature* **620**, 172–180 (2023).
4. Tai-Seale, M. et al. AI-generated draft replies integrated into health records and physicians' electronic communication. *JAMA Netw. Open* **7**, e246565 (2024).
5. Tierney, A. A. et al. Ambient artificial intelligence scribes to alleviate the burden of clinical documentation. *NEJM Catal. Innov. Care Deliv.* **5**, CAT.23.0404 (2024).

6. Jiang, L. Y. et al. Health system-scale language models are all-purpose prediction engines. *Nature* **619**, 357–362 (2023).
7. Cohen, M. K., Kolt, N., Bengio, Y., Hadfield, G. K. & Russell, S. Regulating advanced artificial agents. *Science* **384**, 36–38 (2024).
8. Meskó, B. & Topol, E. J. The imperative for regulatory oversight of large language models (or generative AI) in healthcare. *NPJ Digit. Med.* **6**, 1–6 (2023).
9. Collins, G. S., Reitsma, J. B., Altman, D. G. & Moons, K. G. M. Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): the TRIPOD statement. *BMJ* **350**, g7594 (2015).
10. EQUATOR Network. Reporting guidelines. https://www.equator-network.org/reporting-guidelines/ (accessed 1 June 2024).
11. Collins, G. S. et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ* **385**, e078378 (2024).
12. Liu, X., Cruz Rivera, S., Moher, D., Calvert, M. J. & Denniston, A. K. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat. Med.* **26**, 1364–1374 (2020).
13. Vasey, B. et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nat. Med.* **28**, 924–933 (2022).
14. Norgeot, B. et al. Minimum information about clinical artificial intelligence modeling: the MI-CLAIM checklist. *Nat. Med.* **26**, 1320–1324 (2020).
15. Agrawal, M., Hegselmann, S., Lang, H., Kim, Y. & Sontag, D. Large language models are few-shot clinical information extractors. *Proc. 2022 Conference on Empirical Methods in Natural Language Processing* 1998 (Association for Computational Linguistics, 2022).
16. Liu, X. et al. Large language models are few-shot health learners. Preprint at https://doi.org/10.48550/arXiv.2305.15525 (2023).
17. Jurafsky, D., Chai, J., Schluter, N. & Tetreault, J. (eds). Masked language model scoring. *Proc. 58th Annual Meeting of the Association for Computational Linguistics* 2699–2712 (Association for Computational Linguistics, 2020).
18. Wang, A. et al. GLUE: a multi-task benchmark and analysis platform for natural language understanding. *Proc. 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* 353–355 (Association for Computational Linguistics, 2019).
19. Isabelle, P., Charniak, E. & Lin, D. (eds). Bleu: a method for automatic evaluation of machine translation. *Proc. 40th Annual Meeting of the Association for Computational Linguistics* 311–318 (Association for Computational Linguistics, 2002).
20. Goodman, K. E., Yi, P. H. & Morgan, D. J. AI-generated clinical summaries require more than accuracy. *JAMA* **331**, 637–638 (2024).
21. Zack, T. et al. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *Lancet Digit. Health* **6**, e12–e22 (2024).
22. Gallifant, J. et al. Peer review of GPT-4 technical report and systems card. *PLoS Digit. Health* **3**, e0000417 (2024).
23. Wornow, M. et al. The shaky foundations of large language models and foundation models for electronic health records. *NPJ Digit. Med.* **6**, 135 (2023).
24. Chen, S. et al. The effect of using a large language model to respond to patient messages. *Lancet Digit. Health* **6**, e379–e381 (2024).
25. Chang, C. T. et al. Red teaming large language models in medicine: real-world insights on model behavior. Preprint at *medRxiv* https://doi.org/10.1101/2024.04.05.24305411 (2024).
26. Al-Onaizan, Y., Bansal, M. & Chen, Y.-N. (eds). Language models are surprisingly fragile to drug names in biomedical benchmarks. *Findings of the Association for Computational Linguistics: EMNLP* 12448–12465 (Association for Computational Linguistics, 2024).
27. Boyd, E. Microsoft and Epic expand AI collaboration to accelerate generative AI's impact in healthcare, addressing the industry's most pressing needs. https://blogs.microsoft.com/blog/2023/08/22/microsoft-and-epic-expand-ai-collaboration-to-accelerate-generative-ais-impact-in-healthcare-addressing-the-industrys-most-pressing-needs/ (2023).
28. Moreno, A. C. & Bitterman, D. S. Toward clinical-grade evaluation of large language models. *Int. J. Radiat. Oncol. Biol. Phys.* **118**, 916–920 (2024).
29. Welch Medical Library. Evidence based medicine: evidence grading & reporting. https://browse.welch.jhmi.edu/EBM/EBM_EvidenceGrading (accessed 1 June 2024).
30. Guyatt, G. H. et al. What is 'quality of evidence' and why is it important to clinicians? *BMJ* **336**, 995–998 (2008).
31. Cacciamani, G. E., Collins, G. S. & Gill, I. S. ChatGPT: standard reporting guidelines for responsible use. *Nature* **618**, 238 (2023).
32. Huo, B. et al. Reporting standards for the use of large language model-linked chatbots for health advice. *Nat. Med.* **29**, 2988–2988 (2023).
33. El Mikati, I. K. et al. A framework for the development of living practice guidelines in health care. *Ann. Intern. Med.* **175**, 1154–1160 (2022).
34. Cochrane Community. Living systematic reviews. https://community.cochrane.org/review-development/resources/living-systematic-reviews (accessed 1 June 2024).
35. Akl, E. A. et al. Living systematic reviews: 4. Living guideline recommendations. *J. Clin. Epidemiol.* **91**, 47–53 (2017).
36. Fraile Navarro, D. et al. Methods for living guidelines: early guidance based on practical experience. Paper 5: decisions on methods for evidence synthesis and recommendation development for living guidelines. *J. Clin. Epidemiol.* **155**, 118–128 (2023).
37. Bedoya, A. D. et al. A framework for the oversight and local deployment of safe and high-quality prediction models. *J. Am. Med. Inform. Assoc.* **29**, 1631–1636 (2022).
38. Omiye, J. A., Gui, H., Rezaei, S. J., Zou, J. & Daneshjou, R. Large language models in medicine: the potentials and pitfalls. *Ann. Intern. Med.* **177**, 210–220 (2024).
39. Chen, S. et al. Cross-care: assessing the healthcare implications of pre-training data on language model bias. Preprint at https://doi.org/10.48550/arXiv.2405.05506 (2024).
40. Hansen, L. H. et al. Seeds of stereotypes: a large-scale textual analysis of race and gender associations with diseases in online sources. Preprint at https://doi.org/10.48550/arXiv.2405.05049 (2024).
41. Biderman, S. et al. Pythia: a suite for analyzing large language models across training and scaling. *Proc. 40th International Conference on Machine Learning* 34 (ICML, 2023).
42. Bowman, S. R. et al. Measuring progress on scalable oversight for large language models. Preprint at https://doi.org/10.48550/arXiv.2211.03540 (2022).
43. McAleese, N. et al. LLM critics help catch LLM bugs. Preprint at https://doi.org/10.48550/arXiv.2407.00215 (2024).
44. Burns, C. et al. Weak-to-strong generalization: eliciting strong capabilities with weak supervision. Preprint at https://doi.org/10.48550/arXiv.2312.09390 (2023).
45. Chen, S. et al. Evaluating the ChatGPT family of models for biomedical reasoning and classification. *J. Am. Med. Inform. Assoc.* **31**, 940–948 (2024).
46. Chen, S. et al. Use of artificial intelligence chatbots for cancer treatment information. *JAMA Oncol.* **9**, 1459–1462 (2023).
47. Lu, M. Y. et al. A visual-language foundation model for computational pathology. *Nat. Med.* **30**, 863–874 (2024).
48. Kocak, B. et al. METhodological RadiomICs Score (METRICS): a quality scoring tool for radiomics research endorsed by EuSoMII. *Insights Imaging* **15**, 8 (2024).

49. Lambin, P. et al. Radiomics: the bridge between medical imaging and personalized medicine. *Nat. Rev. Clin. Oncol.* **14**, 749–762 (2017).
50. Shah, N. H. et al. A nationwide network of health AI assurance laboratories. *JAMA* **331**, 245–249 (2024).
51. Diaz, N. Epic releases AI validation suite. *Becker's Hospital Review* https://www.beckershospitalreview.com/ehrs/epic-releases-ai-validation-suite.html (accessed 23 May 2024).
52. diehlbw et al. epic-open-source/seismometer. *GitHub* https://github.com/epic-open-source/seismometer (2024).
53. National Institute of Standards and Technology (NIST). U.S. Artificial Intelligence Safety Institute. https://www.nist.gov/aisi (2023).
54. Federal Register. Health data, technology, and interoperability: certification program updates, algorithm transparency, and information sharing. https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and (2024).
55. EU Artificial Intelligence Act. The AI Act Explorer. https://artificialintelligenceact.eu/ai-act-explorer/ (2024).
56. Peng, C. et al. A study of generative large language model for medical research and healthcare. *NPJ Digit. Med.* **6**, 210 (2023).
57. Zakka, C. Almanac—retrieval-augmented language models for clinical medicine. *NEJM AI* **1**, AIoa2300068 (2024).
58. Keloth, V. K. et al. Advancing entity recognition in biomedicine via instruction tuning of large language models. *Bioinformatics* **40**, btae163 (2024).
59. Guevara, M. et al. Large language models to identify social determinants of health in electronic health records. *NPJ Digit. Med.* **7**, 6 (2024).
60. Jin, Q., Yang, Y., Chen, Q. & Lu, Z. GeneGPT: augmenting large language models with domain tools for improved access to biomedical information. *Bioinformatics* **40**, btae075 (2024).
61. Goh, E. et al. Large language model influence on diagnostic reasoning: a randomized clinical trial. *JAMA Netw. Open* **7**, e2440969 (2024).
62. Zaretsky, J. et al. Generative artificial intelligence to transform inpatient discharge summaries to patient-friendly language and format. *JAMA Netw. Open* **7**, e240357 (2024).
63. Han, L., Erofeev, G., Sorokina, I., Gladkoff, S. & Nenadic, G. Examining large pre-trained language models for machine translation: what you don't know about it. in *Proc. Seventh Conference on Machine Translation (WMT)* https://aclanthology.org/2022.wmt-1.84 (Association for Computational Linguistics, 2022).
64. Yoon, W. et al. LCD benchmark: long clinical document benchmark on mortality prediction. *J. Am. Med. Inform. Assoc.* https://doi.org/10.1093/jamia/ocae287 (2024).

[1]Laboratory for Computational Physiology, Massachusetts Institute of Technology, Cambridge, MA, USA. [2]Department of Critical Care, Guy's and St Thomas' NHS Foundation Trust, London, UK. [3]Artificial Intelligence in Medicine (AIM) Program, Mass General Brigham, Harvard Medical School, Boston, MA, USA. [4]Department of Medicine, University of Wisconsin—Madison, Madison, WI, USA. [5]Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA. [6]Tasmanian School of Medicine, College of Health and Medicine, University of Tasmania, Hobart, Tasmania, Australia. [7]Department of Population Health, NYU Grossman School of Medicine and Langone Health, New York, NY, USA. [8]Department of Radiation Oncology, Brigham and Women's Hospital/Dana–Farber Cancer Institute, Boston, MA, USA. [9]USC Institute of Urology and Catherine and Joseph Aresty Department of Urology, Keck School of Medicine, University of Southern California, Los Angeles, CA, USA. [10]Artificial Intelligence Center, USC Institute of Urology, University of Southern California, Los Angeles, CA, USA. [11]National Library of Medicine, NIH, HHS, Bethesda, MD, USA. [12]Department of Computer Science, Loyola University, Chicago, IL, USA. [13]Department of Dermatology, Stanford School of Medicine, Redwood City, CA, USA. [14]Department of Biomedical Data Science, Stanford School of Medicine, Redwood City, CA, USA. [15]Cognitive Science, Aarhus University, Jens Chr. Skou 2, Aarhus, Denmark. [16]Mass General Brigham, Boston, MA, USA. [17]Faculty of Medicine and Dentistry, University of Alberta, Edmonton, Alberta, Canada. [18]Computational Health Informatics Program, Boston Children's Hospital, Harvard Medical School, Boston, MA, USA. [19]Department of Radiation Oncology, The University of Texas MD Anderson Cancer Center, Houston, TX, USA. [20]Departamento de Telematica, Universidad del Cauca, Popayan, Colombia. [21]Dana–Farber Cancer Institute, Boston, MA, USA. [22]Department of Radiology, Emory University School of Medicine, Atlanta, GA, USA. [23]Centre for Statistics in Medicine, Nuffield Department of Orthopaedics, Rheumatology, and Musculoskeletal Sciences, University of Oxford, Oxford, UK. [24]UK EQUATOR Centre, Nuffield Department of Orthopaedics, Rheumatology, and Musculoskeletal Sciences, University of Oxford, Oxford, UK. [25]Julius Center for Health Sciences and Primary Care, UMC Utrecht, Utrecht University, Utrecht, the Netherlands. [26]Health Innovation Netherlands (HINL), Utrecht, the Netherlands. [27]Division of Pulmonary, Critical Care and Sleep Medicine, Beth Israel Deaconess Medical Center, Boston, MA, USA. [28]Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA. [29]These authors contributed equally: Majid Afshar, Saleem Ameen, Yindalon Aphinyanaphongs, Shan Chen, Giovanni Cacciamani, Dina Demner-Fushman, Dmitriy Dligach, Roxana Daneshjou, Chrystinne Fernandes, Lasse Hyldig Hansen, Adam Landman, Lisa Lehmann, Liam G. McCoy, Timothy Miller, Amy Moreno, Nikolaj Munch, David Restrepo, Guergana Savova, Renato Umeton, Judy Wawira Gichoya.
✉e-mail: dbitterman@bwh.harvard.edu

Methods

The TRIPOD-LLM statement was formulated to guide the reporting of studies that develop, tune or evaluate LLMs for any healthcare application or context and was crafted following pathways used in creating the other TRIPOD statements. An expedited Delphi process was implemented given the need for timely reporting guidelines in this field and was combined with a living statement approach. An accompanying glossary (Box 2) defines essential terms relevant to the TRIPOD-LLM statement.

A steering group was established to direct the guideline development process. They were joined by an expert panel and were chosen based on their diverse expertise and experience in NLP, AI and medical informatics. (The roles of this paper's authors in the two groups are described in Author Contributions.) This guideline was registered on 2 May 2024, with the EQUATOR Network as a reporting guideline under development (https://www.equator-network.org).

Ethics

This study received an exemption from the MIT Committee on the Use of Humans as Experimental Subjects institutional review board (COUHES IRB) on 26 March 2024 (exempt ID: E-5705). Delphi survey participants provided electronic informed consent before completing the survey.

Candidate item list generation

The TRIPOD-2015 and TRIPOD + AI guidelines (https://www.tripod-statement.org) and associated literature on reporting guidelines for LLMs were used to inform the initial candidate item list[9,11,28,32]. The steering group and expert panel expanded this list through additional literature review, ultimately standardizing it to 64 unique items across the following sections: title, abstract, introduction, methods, results, discussion and others.

Panelist recruitment

Delphi participants were identified by the steering committee from authors of relevant publications and through personal recommendations, including experts recommended by other Delphi participants. The steering group identified participants representing geographical and disciplinary diversity, including key stakeholder groups, for example, researchers (statisticians, data scientists, epidemiologists, machine learners, clinicians and ethicists), healthcare professionals, journal editors, funders, policymakers, healthcare regulators and patient advocate groups. No minimum sample size was placed on the number of participants. A steering group member checked the expertise or experience of each identified person. Participants were then invited to complete a survey via email. Delphi participants did not receive any financial incentive or gift to participate.

Delphi process

The survey was designed to allow individual responses in English and delivered electronically using Google Forms. All responses were anonymous; no emails or identifying information was collected from participants. Participants were asked to rate each item as 'can be omitted', 'possibly include', 'desirable for inclusion' or 'essential for inclusion', as has been conducted in previous TRIPOD guidelines[9]. Participants were also invited to comment on any item and suggest new items. D.S.B. and J.G. collated and analyzed the free-text responses and then used the themes generated to inform item rephrasing, merging or suggesting new items. All steering group members were invited to participate in the Delphi surveys.

Round 1 participants

The first round was conducted between 1 March 2024 and 23 April 2024, where the participation link was sent to 56 people. Of the 56 invited, 26 completed the survey. Survey participants came from nine countries, with 14 from North America, 5 from Europe, 2 from Asia, 1 from South America and 1 from Australasia. Three participants did not provide this information. Participants reported their primary fields of work and could select more than one field. In total, 20 of 26 (77%) had a primary field in AI, machine learning, clinical informatics or NLP, and 14 of 26 selected healthcare.

Consensus meeting

An online consensus meeting was held on the 22nd and 24th of April, chaired by D.S.B. and J.G. via Zoom. All steering committee and expert panel members were invited to attend the meeting. Recordings and

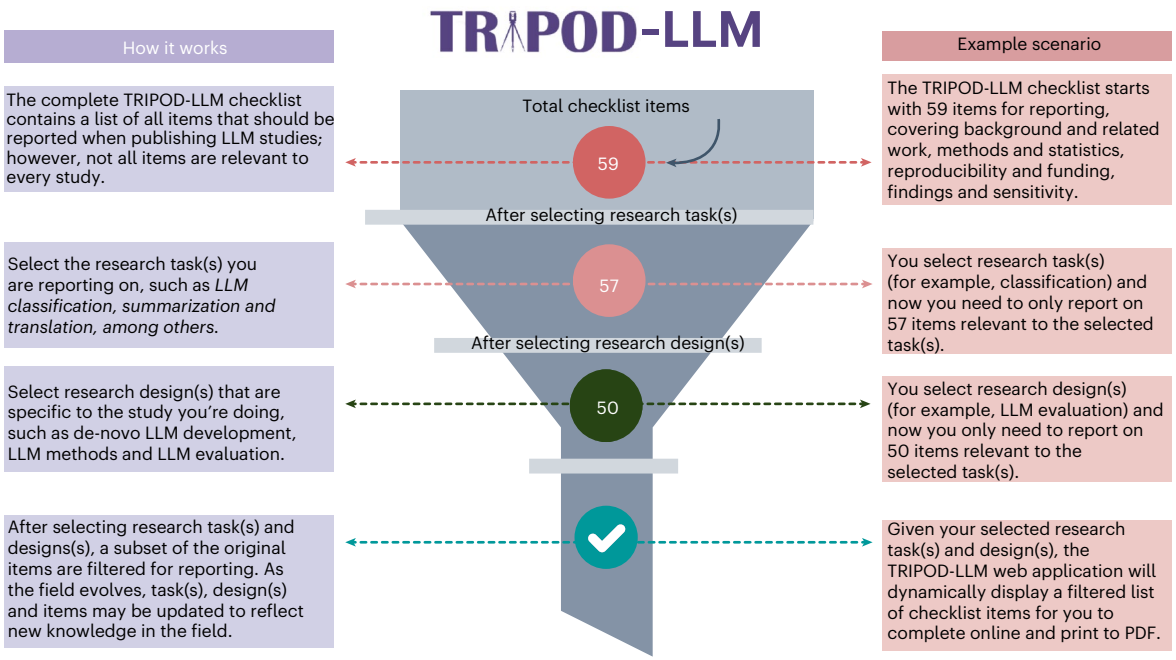

Fig. 1 | TRIPOD-LLM workflow. The TRIPOD-LLM checklist workflow starts with 59 reporting items, and the number of required items is reduced based on the selection of research tasks (for example, classification and summarization) and research design (for example, LLM evaluation). After selecting both, a filtered list is generated for reporting.

notes were sent immediately after the meetings to enable asynchronous contribution for those who did not attend. The responses to each question were examined in turn, as well as all free-text comments. Items with support from <50% of participants for being 'Essential to Include' were highlighted and deliberated for the importance of inclusion. Agreement by consensus was reached in all cases, without needing a third party. To arrive at a consensus, the item was discussed until no panel member had additional comments or disagreements with the final disposition of the item. Transcripts of the discussion, which have been de-identified and any sensitive information removed, are available in Supplementary Information for full transparency.

Due to the vast number of applications being developed using LLMs, a modular approach was used to group included items under additional subcategories under the 'Research Design' and 'LLM Task' headers. This approach was agreed upon during the meeting, and the steering committee approved the final groupings.

To adequately address the variety of studies and uses regarding LLMs, ranging across stages of development, tuning, evaluation and implementation, as well as across healthcare tasks, items are categorized according to (1) research design and (2) LLM task (Fig. 1). The research design categories are de novo LLM development, LLM methods such as fine-tuning, prompt-engineering techniques and architecture modifications, intrinsic LLM evaluation and LLM evaluation in dedicated healthcare settings and tasks. The LLM task categories are lower-level text processing (for example, part-of-speech tagging, relation extraction and named entity recognition), classification (for example, diagnosis), long-form question answering, conversational agent, documentation generation, summarization/simplification, machine translation and outcome forecasting (for example, prognosis). Items may apply to several design and task categories, and studies may include more than one design and task. Items applicable to every type and task covered in the study should be reported. Definitions and examples of each design and task category are provided in Table 1. We acknowledge that these categorizations are imperfect, and overlap may exist across designs and tasks.

Acknowledgements

J.G. is funded by the US National Institutes of Health (NIH) (U54 TW012043-01 and OT2OD032701). S.C. is supported by the NIH (U54CA274516-01A1). G.S.C. is supported by Cancer Research UK (program grant C49297/A27294) and by the Engineering and Physical Sciences Research Council grant for 'Artificial intelligence innovation to accelerate health research' (EP/Y018516/1), and is a National Institute for Health and Care Research (NIHR) senior investigator. The views expressed in this article are those of the author and not necessarily those of the NIHR or the Department of Health and Social Care. Y.A. is partially supported by the NIH (3UL1TR001445-05) and the National Science Foundation Award (1928614 and 2129076). L.A.C. is funded by the NIH (U54 TW012043-01, OT2OD032701 and R01EB017205). D.S.B. is supported by the NIH (U54CA274516-01A1 and R01CA294033-01) and the American Cancer Society and American Society for Radiation Oncology grant ASTRO-CSDG-24-1244514-01-CTPS. J.W.G. receives support from 2022 Robert Wood Johnson Foundation Harold Amos Medical Faculty Development Program, a Radiological Society of North America Health Disparities grant (EIHD2204), the Lacuna Fund (67), the Gordon and Betty Moore Foundation, an NIH (National Institute of Biomedical Imaging and Bioengineering) MIDRC grant under contracts 75N92020C00008 and 75N92020C00021, and an NHLBI award (R01HL167811).

Author contributions

D.S.B., J.G., L.A.C., G.S.C. and K.G.M.M. were on the steering group that directed the guideline-development process. S.C., C.F., D.R., G.S., T.M., D.D.F., R.U., L.H.H., Y.A., J.W.G., L.G.M., N.M. and R.D. were members of the expert panel. D.S.B. and J.G. drafted the initial list of candidate items.

Competing interests

D.S.B. is an associate editor at *Radiation Oncology* and HemOnc.org, receives research funding from the American Association for Cancer Research, and provides advisory and consulting services for MercurialAI. D.D.F. is an associate editor at the *Journal of the American Medical Informatics Association*, is a member of the editorial board of *Scientific Data*, and receives funding from the intramural research program at the US National Library of Medicine, NIH. J.W.G. is a member of the editorial board of *Radiology: Artificial Intelligence*, *BJR Artificial Intelligence* and *NEJM AI*. All other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41591-024-03425-5.

Correspondence and requests for materials should be addressed to Danielle S. Bitterman.

Peer review information *Nature Medicine* thanks Michael Sorich, Fares Antaki and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editor: Karen O'Leary, in collaboration with the *Nature Medicine* team.