

Towards AI-Assisted Clinical Trial Matching: Practical Considerations, Multicenter Evaluation, and Real-World Deployment

Yin Fang¹, Qiao Jin¹, Shubo Tian¹, Lauren He¹, Maya Geer¹, Noor Naffakh², Ryan Huu-Tuan Nguyen², Zifeng Wang³, Jimeng Sun³, Charalampos S. Floudas⁴, James L. Gulley⁴, Kamilia Moalem⁵, Catarina Martins Maia⁵, Amanda Nottke⁶, Juan W. Valle^{6,9}, Melinda Bachini⁶, Lourdes Rocha-Nussbaum⁶, Kari Ramage⁶, Nikita Curry⁷, Megan Barnes⁷, Mandy Mansaray⁷, Darlene Gabeau⁸, Craig E. Grossman⁸, Heath Skinner⁸, Michael Burczynski⁸, NIH-TrialBench Consortium, Zhiyong Lu^{1,*}

¹Division of Intramural Research, National Library of Medicine, National Institutes of Health, Bethesda, MD, USA.

²UI Health, University of Illinois Chicago. Chicago, IL, USA.

³University of Illinois Urbana-Champaign, Urbana, IL, USA.

⁴Center for Immuno-Oncology, Center for Cancer Research, National Cancer Institute, National Institutes of Health, Bethesda, MD, USA.

⁵Medical Oncology Service, Center for Cancer Research, National Cancer Institute, National Institutes of Health, Bethesda, MD, USA

⁶The Cholangiocarcinoma Foundation, Herriman, UT, USA.

⁷Office of Patient Recruitment, NIH Clinical Center, Bethesda, MD, USA.

⁸UPMC Hillman Cancer Center, University of Pittsburgh, Pittsburgh, PA, USA.

⁹Division of Cancer Sciences, University of Manchester, Manchester, UK.

*Corresponding author: zhiyong.lu@nih.gov

Abstract

Clinical trials are essential for advancing cancer care and drug development, but many fail because of insufficient patient enrollment. While there is growing interest in using AI to support patient recruitment, existing systems largely perform eligibility assessment alone and have rarely been evaluated in real-world oncology workflows. Here we present TrialGPT 2.0, an AI-assisted clinical trial recommendation system designed for real-world deployment. Rather than asking only whether a patient may qualify, the system also assesses which trials warrant further consideration given the patient's current clinical needs and local workflow priorities, and provides structured, inspectable explanations for expert review. Importantly, we evaluated TrialGPT 2.0 retrospectively and prospectively across multiple oncology-focused settings, spanning government, academic cancer-center, patient-advocacy, and NIH referral workflows. In retrospective multicenter cohorts comprising 288 cases, TrialGPT 2.0 retrieved at least one clinician-recommended trial in its top 10 recommendations for approximately 91% of cases while reducing clinician screening time by 55.0%. In a six-month prospective evaluation embedded in an active precision oncology tumor board, TrialGPT 2.0 contributed additional trial opportunities missed by the routine workflow, expanding patient access to clinical trial participation by 90.9%. To support scientific reproducibility, we also introduce NIH-TrialBench, a clinician-authored dataset comprising 126 diverse synthetic patient vignettes and matching scenarios from 11 NIH Institutes and Centers. Together, these results support the value of AI to assist clinical trial matching by improving clinician efficiency and identifying frequently overlooked trial opportunities, ultimately helping to expand and accelerate accrual to cancer trials. TrialGPT 2.0 has been deployed in real-world trial-review workflows.

Introduction

Clinical trials provide evidence that guides medical practice and remain the gold standard for drug development, with a particularly important role in advancing cancer care. However, bringing a new drug to market is resource-intensive, often requiring over a decade and more than a billion dollars, with clinical trials accounting for over half of these expenditures^{1,2}. Patient-trial matching is a central operational step in recruitment, particularly in oncology, where it requires interpretation of complex and rapidly evolving clinical and molecular eligibility criteria, integration of heterogeneous patient information for each candidate trial, and continuous review of the changing portfolio of trials available at a given site³. Today, this step remains largely manual, making it slow and error-prone⁴. Poor patient recruitment is a leading driver of research waste in randomized controlled trials (RCTs), accounting for 36.7% of premature discontinuations, which affect 30.1% of general RCTs⁵. The challenge is especially acute in oncology: 22.7% of interventional cancer trials terminate early, with accrual problems responsible for 34.5% of those terminations^{6,7}, and low accrual accounts for 57.0% of failed Phase II/III cancer trials^{8,9}.

To address this bottleneck, various artificial intelligence (AI) methods have been explored to support patient recruitment through trial search^{10–13}, cohort identification^{14,15} and eligibility screening^{16–25}. For example, we previously developed TrialGPT, a first-of-its-kind large language model (LLM)-based framework for zero-shot patient-to-trial matching (i.e., without task-specific training) that combines retrieval, criterion-level eligibility assessment, and trial-level ranking, into a single system¹⁰.

Despite this progress, existing AI-assisted patient-trial matching approaches share two major limitations that constrain their clinical value. First, they remain focused on trial eligibility rather than clinical recommendation. These are distinct questions: eligibility establishes whether a patient could plausibly enroll, whereas recommendation reflects whether the trial is appropriate to pursue given the patient’s current clinical needs. This distinction is especially important in oncology, where a patient may be eligible for many trials but warrant recommendation for only a few (e.g., a clinician might not recommend an eligible trial because it is purely diagnostic, offers little expected benefit relative to standard care, or contradicts the patient’s overall treatment plan). Therefore, eligibility is necessary but not sufficient for recommendation. Second, prior evaluations have relied largely on technical validation using synthetic or eligibility-oriented benchmarks and have provided limited evidence from clinician-adjudicated or prospective oncology review workflows with real patient data, leaving it unclear whether AI-identified trials contribute useful options to real-world trial-selection decisions^{17,18}. Together, these gaps point to an urgent and unmet need in medical AI research: to evaluate patient-trial matching as a clinical recommendation task aligned with routine clinical workflows, rather than on benchmark performance alone^{26–28}.

To this end, we present TrialGPT 2.0, an enhanced TrialGPT¹⁰ system that designed to assist clinicians and trial specialists in real-world patient-trial matching. Unlike its predecessor, TrialGPT 2.0 is designed with greater configurability to support the trial-review needs of different recruitment settings and workflows. Each setting uses a predefined trial list relevant to its workflow (e.g., an institutional or program-specific portfolio) and a matching policy that specifies how potentially eligible trials should be prioritized (e.g., by favoring stronger disease or biomarker matches or therapeutic rather than diagnostic or registry trials). Given a patient, TrialGPT 2.0 generates a ranked list of candidate trials based on overall patient-trial fit. Each result includes a structured assessment separating eligibility-supporting evidence, potential incompatibilities, and missing information, alongside a rationale to support transparent clinical review.

More importantly, we evaluated TrialGPT 2.0 across complementary retrospective and prospective cohorts, predominantly in oncology. First, we performed a multicenter retrospective evaluation of 1,340 heterogeneous patient-trial pairs from 288 de-identified clinical-note cases from four oncology-specific workflows and one broader NIH referral workflow. These cases covered referral, intake, consultation, and tumor-board workflows and varied in patient characteristics, documentation length, and local trial search space. Second, we prospectively assessed TrialGPT 2.0 in a Precision Oncology Tumor Board (POTB) trial-review task²⁹. From February through July 2026, TrialGPT 2.0 outputs were reviewed during routine clinical practice across 27 cases and 339 patient-trial pairs to assess whether the system contributed final trial options beyond the conventional human workflow. These evaluations directly address the two limitations identified above: they assess recommendation quality beyond eligibility through domain-expert clinical review, and they test whether AI-identified options are selected into final recommendations in a real-world oncology workflow.

Additionally, to support reproducible evaluation despite restrictions on sharing real clinical data, we developed NIH-TrialBench, a clinician-authored benchmark comprising 126 synthetic vignettes spanning oncology and other disease areas and contributed by investigators from 11 NIH Institutes and Centers. We evaluated TrialGPT 2.0 on both NIH-TrialBench and established public patient-trial matching benchmarks^{30–32}.

Finally, we implemented and deployed these capabilities in a web-based interface for routine trial review in real-world clinical workflows. Together, these results support AI-assisted matching as a practical means of helping oncology teams identify clinically appropriate trial options more efficiently, an important step toward broadening and accelerating cancer trial recruitment.

Results

TrialGPT 2.0 provides configurable and explainable clinical trial recommendations

TrialGPT 2.0 is an AI-powered clinical trial matching system that assists clinicians in identifying and recommending clinical trials for individual patients. Here, recommendation refers not only to apparent eligibility but also to whether a trial warrants further consideration given the patient's current clinical needs. For input, the system takes patient information, a predefined list of available trials relevant to the recruitment setting (e.g., an institutional portfolio or a list identified through a disease-focused or geographic query), and a local matching policy that reflects the clinical preferences (e.g., disease- or biomarker-specific priorities, or a preference for therapeutic over diagnostic or registry trials) for prioritizing eligible trials (Fig. 1a). For output, TrialGPT 2.0 generates a ranked list of trial recommendations for clinical review (Fig. 1c).

To accommodate varying numbers of applicable trials at different sites and workflows, TrialGPT 2.0 applies a pluggable retrieval component before detailed matching assessment. Across the study cohorts, the local trial list ranged from 9 to 1,871 candidate trials, scoped by disease or biomarker focus, geographic region or institutional availability (Fig. 2a). For local lists exceeding 1,500 trials, the retrieval module is automatically activated to generate a candidate shortlist for improved efficiency, using the previously described hybrid-fusion strategy of TrialGPT that combines lexical and semantic retrieval¹⁰; otherwise, the full list of trials is passed directly to the backend. The backend then generates an assessment for each candidate trial by comparing the patient with the trial's own inclusion and exclusion criteria and evaluating relevant trial-fit factors defined by the matching policy for the recruitment setting. These factors may include diagnosis




Fig. 1. Overview of the TrialGPT 2.0 system and evaluation framework. **a**, Inputs for patient-trial matching, including the local trial corpus, patient context and center-specific matching policies. **b**, TrialGPT 2.0 system architecture. A pluggable retrieval module supports trial corpora of different sizes, and the backend evaluates trial-fit factors to generate structured trial-level assessments, including eligibility evidence, missing information, rationale, fit score, confidence, recommendation and rank. **c**, Ranked output interface with post-ranking filters, sorting options, trial-level explanations, scores, confidence estimates and recommendation categories. **d**, Evaluation framework. TrialGPT 2.0 is evaluated through retrospective clinician-adjudicated review, prospective workflow evaluation, reproducible benchmarking, and real-world deployment evaluation.

or histology, disease status, biomarker or pathway match, line of therapy, prior treatment, and cohort or arm applicability. Each matching policy reflects the clinical priorities and information available in the corresponding recruitment setting by specifying which factors are explicitly assessed and which are considered only when relevant information is available (Supplementary Note 1). Subsequently, for each candidate trial, TrialGPT 2.0 returns a recommendation category (Highly Recommended, Possible Match, or Low Fit), a rank, a fit score, a confidence estimate, and a structured explanation separating eligible versus ineligible reasons, missing information, and rationale, making both the recommendation and its supporting evidence inspectable (Fig. 1b). To facilitate review, we implemented a web-based interface that presents ranked trial cards with these structured explanations and supports customizable post-ranking filtering and sorting aligned with local workflow constraints and referral priorities (Supplementary Fig. 4).

We evaluated TrialGPT 2.0 through four complementary components (Fig. 1d), with the clinical evaluation spanning four oncology-specific trial-selection workflows and one broader NIH referral workflow across government, patient-advocacy, and academic cancer-center settings (Fig. 2a).

First, we performed a retrospective, clinician-adjudicated review of recommendation quality across de-identified clinical-note workflows. This evaluation comprised 288 de-identified cases from five




Fig. 2. Characteristics of study cohorts. **a**, Summary of study cohorts. “Top-ranked review” indicates clinician labeling of up to 10 top-ranked TrialGPT 2.0 Highly Recommended trials per case. “Exhaustive review” indicates clinician labeling of all trials in the search space. “Prospective review” indicates workflow-based review in the prospective POTB setting. **b**, Disease-area coverage across cohorts. Patient contexts are mapped to first-level MeSH³³ disease categories, with individual cases allowed to span multiple disease areas. Horizontal bars show the percentage of disease-area assignments in each MeSH category, and dots indicate whether each category was represented in each cohort type. **c**, Age and sex distributions across study cohorts. Horizontal bars show age-group counts, and overlaid pie markers indicate sex composition within each age group.

recruitment pathways: patient call intake notes from the NIH Office of Patient Recruitment (OPR); oncology patient advocacy intake notes from the Cholangiocarcinoma Foundation (CCF); tumor-board referrals from the University of Illinois Cancer Center (UIC); radiation-oncology consultation notes from a community cancer program within the University of Pittsburgh Medical Center (UPMC) Hillman Cancer Center; and oncology referral notes from the National Cancer Institute (NCI). These pathways represented distinct matching needs: broad referral triage; disease- and biomarker-specific navigation for cholangiocarcinoma and biliary tract cancer; molecularly guided trial prioritization in precision oncology; radiation-oncology treatment-fit assessment; and comprehensive referral screening within a focused institutional oncology portfolio. The patient documents differed substantially in format and detail, ranging from short OPR and CCF intake or advocacy summaries and concise UIC molecular tumor-board referrals to curated NCI referral summaries and longer UPMC narrative or document-derived clinical notes. Each site defined its own candidate trial list using a workflow-appropriate scoping strategy. CCF and UIC derived their lists from ClinicalTrials.gov³⁴ using disease- and biomarker-specific and location-based queries, respectively; OPR used the NIH Clinical Center protocol inventory³⁵, NCI used a focused institutional oncology set, and UPMC used a manually curated radiation-oncology portfolio. Matching was performed independently for each setting against its predefined list, reflecting how trial review is scoped in practice, where clinical teams generally consider a relevant institutional, geographic, disease-focused, or program-specific trial set rather than the full registry. Accordingly, most search spaces contained fewer than 1,500 candidate trials (Fig. 2a).

Second, we prospectively evaluated whether TrialGPT 2.0-identified options contributed to final recommendations in the University of Illinois Cancer Center Precision Oncology Tumor Board (UIC POTB), a routine multidisciplinary oncology workflow for reviewing referred cases and identifying clinical trial options. The evaluation included 27 de-identified cases. For each case, TrialGPT 2.0 and clinical reviewers identified candidate trials in parallel from the same patient context. The two candidate-trial lists were then reviewed together, and the board selected the final trial recommendations. TrialGPT 2.0 used the UIC trial-corpus snapshot available at the time of each case review, which was refreshed weekly to reflect changes in trial availability, recruitment status, and cohort or arm availability (Fig. 2a).

Third, we conducted a technical validation on existing eligibility-focused public patient-trial matching benchmarks and on NIH-TrialBench, a new benchmark we developed comprising 126 synthetic vignettes spanning oncology and other disease areas. The vignettes were created by 24 NIH trial investigators across 11 NIH Institutes and Centers to reflect real-world considerations in trial enrollment and screening. These vignettes were evaluated against a predefined search space of 1,373 candidate trials from the NIH Clinical Center's *Search the Studies* database³⁵, and clinicians reviewed and labeled the top-ranked TrialGPT 2.0 recommendations for each vignette. Unlike the data in the first two evaluations, NIH-TrialBench uses synthetic vignettes to support reproducible evaluation without exposing private patient information (Fig. 2a).

Fourth, we evaluated real-world deployment through a pilot user survey of clinical reviewer acceptance and interface runtime testing (Fig. 1d).

These components provided complementary coverage of evaluation settings, disease areas, patient demographics, and documentation patterns. Neoplasms were the largest disease-area category, reflecting the predominant oncology focus of the clinical workflows, while the combined evaluation also spanned a broad range of Medical Subject Headings (MeSH)³³ categories. Retrospective clinical notes and clinician-authored synthetic vignettes showed the broadest disease-area coverage, whereas prospective cases were concentrated in the precision oncology setting (Fig. 2b).

Age and sex distributions also differed across cohort types (Fig. 2c).

TrialGPT 2.0 returns clinician-recommended trials and improves screening efficiency in multicenter retrospective evaluations

We evaluated TrialGPT 2.0 in retrospective clinical-note cohorts from four oncology-specific recruitment settings and one broader NIH referral setting across the United States (Fig. 3a) to assess whether its recommendations aligned with clinician judgment and whether TrialGPT 2.0 assistance reduced clinician screening time.

We first assessed top-ranked recommendations in the OPR, CCF, UIC, and UPMC cohorts, spanning broad NIH referral triage, cholangiocarcinoma trial navigation, precision-oncology tumor-board review, and radiation-oncology consultation. Clinicians reviewed up to ten top-ranked trials categorized by TrialGPT 2.0 as Highly Recommended (or all such trials if the system identified fewer than ten), and labeled each patient-trial pair as Recommended, Eligible but Not Recommended, or Ineligible. We summarized performance using hit rate@ K (the fraction of cases with at least one trial judged Recommended in the top K), recommended precision@ K (the fraction of top- K recommendations judged Recommended), and eligible precision@ K (the fraction of top- K recommendations judged Recommended or Eligible but Not Recommended). We designated hit rate@ K the primary metric because it best reflects the intended workflow, in which a short ranked list should contain at least one clinician-recommended option.

Across these four cohorts, clinicians reviewed 440 patient-trial pairs from 188 de-identified cases spanning short call intake notes, brief advocacy summaries, concise tumor-board referrals, and longer consultation notes. Hit rate was already 0.88 at $K = 1$, indicating that the top-ranked recommendation was clinician-recommended in 88% of cases. Hit rate reached 0.91 at $K = 3$ and remained at 0.91 through $K = 10$, indicating that the evaluated top-10 list contained at least one clinician-recommended trial in 91% of cases. Eligible precision remained high across cutoffs, decreasing only slightly from 0.89 at $K = 1$ to 0.87 at $K = 10$. Recommended precision, which required clinicians to judge a trial appropriate to pursue rather than merely eligible, was 0.88 at $K = 1$ and 0.83 at $K = 10$ (Fig. 3b).

For the National Cancer Institute (NCI) cohort, comprising de-identified oncology referral notes from the Center for Immuno-Oncology (CIO) of the Center for Cancer Research (CCR), the predefined search space comprised a focused institutional portfolio of nine oncology trials, enabling exhaustive review of all nine trials for each of the 100 cases and yielding 900 patient-trial pairs. For the agreement analysis, the original clinician labels were mapped to the three TrialGPT 2.0 output categories as described in Methods. Each pair received three human annotations: one from an independent full-cohort review and two from a counterbalanced experiment in which two clinicians each reviewed all pairs, with TrialGPT 2.0 assistance assigned to opposite 50-case halves of the cohort. The resulting clinician consensus labels served as the reference standard.

Against these reference labels, TrialGPT 2.0 alone reached 89.0% exact agreement, whereas clinician review with and without TrialGPT 2.0 reached 94.0% and 94.2%, respectively. TrialGPT 2.0 correctly identified 103 of 118 Highly Recommended pairs and 671 of 723 Low Fit pairs, corresponding to class-specific recalls of 87.3% and 92.8%, respectively. Agreement was lower for Possible Match pairs (27 of 59), consistent with the greater ambiguity of intermediate cases, for which incomplete information can affect adjudication. Although clinician review with and without assistance achieved similar overall agreement, the error profiles differed. Severe over-recommendation, defined as labeling a consensus-reference Low Fit pair as Highly Recommended,




Fig. 3. Retrospective evaluation of TrialGPT 2.0 recommendations and clinician screening time. **a**, Geographic locations of the five participating retrospective cohort sites, comprising four oncology-specific trial-selection settings and one broader NIH referral workflow. **b**, Recommendation quality performance on retrospective clinical-note cohorts from OPR, CCF, UIC, and UPMC. Bars show hit rate@K, recommended precision@K, and eligible precision@K for $K = 1, 3, 5$ and 10, with error bars showing 95% confidence intervals of the mean computed by percentile bootstrap (10,000 resamples). **c**, Agreement with consensus reference labels in the NCI exhaustive review setting. Red boxes mark severe reversals between Highly Recommended and Low Fit. **d**, Total trial-level screening decision time per case in the NCI counterbalanced experiment. Points represent cases, lines connect paired assisted and unassisted reviews, and brackets show percentage reductions. Bars show means and error bars show 95% bootstrap confidence intervals (10,000 resamples). Statistical comparisons used two-sided Wilcoxon signed-rank tests for paired reviews and a two-sided Mann-Whitney U test between clinicians. All 100 cases were analyzed; one outlier falls outside the displayed range.

occurred in 1 pair with TrialGPT 2.0-assisted review, compared with 14 without assistance and 8 for TrialGPT 2.0 alone. Across both directions, assisted review produced 5 severe reversals between Highly Recommended and Low Fit, compared with 16 without assistance and 11 for TrialGPT 2.0 alone. This pattern suggests that TrialGPT 2.0 assistance may reduce clinically concerning reversals while preserving high agreement with consensus labels (Fig. 3c).

Finally, for the NCI/CCR/CIO cohort, we assessed clinician screening time in the same counter-balanced experiment, defined as clinician-recorded trial-level decision time summed across the nine candidate trials (excluding patient-context review and automated TrialGPT 2.0 PDF report generation). TrialGPT 2.0 assistance reduced mean screening time per patient from 129.8 to 58.4 seconds, a 55.0% reduction, with consistent effects in both groups (100.9 to 41.2 seconds; 158.8 to 75.5 seconds). Mean screening time did not differ significantly between the two clinicians (88.2 versus 100.0 seconds per patient), indicating that the reduction was not explained by a clinician-level speed difference. Overall, these results indicate that TrialGPT 2.0 can identify clinician-recommended options across diverse oncology and referral workflows, improve screening efficiency, and maintain high agreement with expert consensus labels (Fig. 3d).

TrialGPT 2.0 substantially expands clinician-selected trial recommendations in prospective precision oncology review

We prospectively evaluated TrialGPT 2.0 in the UIC POTB, a routine multidisciplinary precision-oncology workflow that reviews clinical context, molecular findings, prior therapies, and the clinical question to identify relevant trial options. Cases entered through routine POTB referral. Trial identification proceeded in parallel: clinical reviewers identified candidate trials through the routine workflow, while TrialGPT 2.0 generated recommendations from the same case context. The two lists were reviewed together, and clinicians selected the final POTB recommendation list. The primary endpoint was TrialGPT 2.0's contribution to the final clinician-selected recommendations, assessed at the case level (whether TrialGPT 2.0 contributed at least one final clinician-selected recommendation in cases where the routine workflow found none) and at the recommendation level (the source of each final recommendation) (Fig. 4a,b). The secondary endpoint characterized all TrialGPT 2.0-generated candidate trials using clinician labels (Fig. 4c).

At the case level, the analysis included 27 POTB cases reviewed from February through July 2026 (Fig. 4a). TrialGPT 2.0 identified final recommendations in 10 cases for which the routine workflow had found none, increasing the number of cases with at least one final recommendation from 11 to 21, a 90.9% increase. Considering all 27 reviewed cases, TrialGPT 2.0 expanded final trial options in 10 cases (37.0%), left opportunities unchanged in 11 (40.7%) where the routine workflow had already identified at least one recommendation, and neither source contributed a final recommendation in 6 cases (22.2%).

At the recommendation level, clinicians selected 54 final trial recommendations, of which TrialGPT 2.0 identified 45 (83.3%): 37 (68.5%) by TrialGPT 2.0 alone and 8 (14.8%) by both TrialGPT 2.0 and the routine workflow (Fig. 4b). The routine workflow identified 17 recommendations in total (the 8 shared with TrialGPT 2.0 and 9 identified by the routine workflow alone). The 37 recommendations uniquely contributed by TrialGPT 2.0 therefore more than doubled the final recommendation set, expanding it from 17 to 54.

For the secondary endpoint, we analyzed clinician labels for the 339 candidate trials TrialGPT 2.0 generated across the 27 cases (Fig. 4c). UIC clinicians labeled 45 (13.3%) as Recommended, 282 (83.2%) as Eligible but Not Recommended, and 12 (3.5%) as Ineligible. The predominance




Fig. 4. Prospective evaluation of TrialGPT 2.0 in Precision Oncology Tumor Board (POTB) trial review. **a**, Case-level source attribution of final clinician-selected trial recommendations across 27 POTB cases reviewed from February through July 2026. Bars show the number of final recommendations identified by TrialGPT 2.0 only, by both TrialGPT 2.0 and the routine workflow, or by the routine workflow only; cases without a final recommendation are shown as zero. The donut chart summarizes TrialGPT 2.0's case-level contribution: cases in which TrialGPT 2.0 expanded trial opportunities by contributing at least one final recommendation where the routine workflow identified none; cases in which opportunities were unchanged because the routine workflow had already identified at least one final recommendation; and cases with no final recommendation from either source. Stars mark the cases expanded by TrialGPT 2.0. **b**, Monthly and overall source attribution of the 54 final clinician-selected trial recommendations. Counts and percentages indicate recommendations identified by TrialGPT 2.0 only, by both TrialGPT 2.0 and the routine workflow, or by the routine workflow only. **c**, Monthly and overall clinician adjudication of the 339 candidate trials generated by TrialGPT 2.0. Candidate trials were labeled as Recommended, Eligible but Not Recommended or Ineligible. Recommended indicates selection for the final POTB recommendation list; Eligible but Not Recommended indicates a trial considered eligible but not prioritized for final recommendation; and Ineligible indicates incompatibility with the patient context or trial criteria.

of Eligible but Not Recommended labels directly illustrates why eligibility alone is insufficient in precision oncology: most candidate trials were not clearly incompatible, yet only a small subset was prioritized for the final recommendation list. Among high-scoring candidates labeled Eligible but Not Recommended, clinicians deprioritized technically eligible trials when poor performance status or comorbidities made intensive regimens less appropriate, when standard-of-care or more specific targeted options were preferred, when biomarker evidence was limited, or when the POTB question was diagnostic rather than therapeutic. These patterns indicate that clinical recommendation in precision oncology requires expert integration of treatment intent, expected benefit, molecular relevance, evidence strength, feasibility, and patient-specific context beyond protocol eligibility alone.

As an exploratory failure-mode analysis, we reviewed the three cases in which all final recommendations came from the routine workflow alone. In each, TrialGPT 2.0 applied protocol eligibility criteria literally, whereas clinicians interpreted the same information using broader oncologic context. In one case, TrialGPT 2.0 excluded a patient from a first-line trial because the record showed prior chemotherapy; clinicians recognized that this therapy had been given when the patient’s lung lesions were initially misclassified as metastases from another primary, leaving the lung tumor itself effectively untreated. In another, TrialGPT 2.0 rejected a strong biomarker-matched trial because the protocol required progression on a specific prior drug with no intervening treatment, whereas the patient had received an interim therapy. In the third, TrialGPT 2.0 down-ranked biologically relevant trials because poor performance status and severe comorbidities, including advanced dementia, raised eligibility concerns. In each case, clinicians drew on specialist judgment about tumor biology, attribution of prior treatment, and trial feasibility to retain these trials as reasonable options for discussion.

NIH-TrialBench supports reproducible benchmarking for AI-assisted clinical trial matching

To complement the retrospective and prospective evaluations, we assessed TrialGPT 2.0 in two reproducible benchmarking settings: NIH-TrialBench and established public benchmarks. Unlike existing eligibility-oriented public benchmarks, NIH-TrialBench evaluates clinician-adjudicated labels against a realistic, workflow-defined trial search space and includes a prespecified target trial for each vignette, enabling evaluation of clinical recommendation quality and target-trial recovery (Fig. 5a).


On NIH-TrialBench, using the same top-ranked review framework as the retrospective cohorts (Fig. 3b), clinicians reviewed up to ten top-ranked trials categorized by TrialGPT 2.0 as Highly Recommended for each vignette and labeled each patient-trial pair as Recommended, Eligible but Not Recommended, or Ineligible. Across 126 vignettes and 990 reviewed patient-trial pairs, eligible precision was 0.92 at $K = 1$ and 0.86 at $K = 10$, and hit rate increased from 0.63 at $K = 1$ to 0.95 at $K = 10$, indicating that the top-10 recommendation list included at least one clinician-recommended trial for most vignettes (Fig. 5b).

Because each vignette was constructed around a known intended target trial, NIH-TrialBench also enabled direct assessment of target-trial recovery, measured as Recall@10 (the proportion of vignettes for which the known target trial appeared among the ten highest-ranked trials). TrialGPT 2.0 achieved a macro-averaged target-trial Recall@10 of 70% across the 11 contributing NIH Institutes and Centers, compared with 54% for the original TrialGPT framework¹⁰ (hereafter TrialGPT 1.0) and, among off-the-shelf general-purpose models, 43% for ChatGPT-5.5 Thinking³⁶ and 21% for Gemini-3.5 Flash Extended Thinking³⁷, indicating stronger recovery of intended


a

Benchmark	# Cases	Mean length (words)	Realistic trial search space	Recommendation vs. eligibility	Prespecified target trial
SIGIR 2016	58	89	No	Eligibility	No
TREC 2021	75	156	No	Eligibility	No
TREC 2022	50	110	No	Eligibility	No
NIH-TrialBench	126	116	Yes	Recommendation	Yes

b



c



d




Fig. 5. TrialGPT 2.0 performance in reproducible benchmarking settings. a, Comparison of NIH-TrialBench with established public benchmarks across characteristics. **b,** NIH-TrialBench top-ranked review performance. Bars show hit rate@K, recommended precision@K, and eligible precision@K for $K = 1, 3, 5$ and 10 , with error bars showing 95% confidence intervals of the mean computed by percentile bootstrap (10,000 resamples). The heatmap summarizes the number of vignettes, reviewed patient-trial pairs, clinician labels and average vignette length for each contributing NIH Institute or Center and overall. **c,** Target-trial Recall@10 on NIH-TrialBench, comparing TrialGPT 2.0 with TrialGPT 1.0 and off-the-shelf general-purpose models, shown per contributing NIH Institute or Center and as the mean across institutes. Error bars show 95% confidence intervals computed by percentile bootstrap (10,000 resamples), and annotations indicate the percentage-point difference between TrialGPT 2.0 and each comparator model. **d,** Public benchmark comparison. TrialGPT 2.0 is compared with TrialGPT 1.0 using Precision@10, nDCG@10, MRR and MAP. Bars show mean performance, error bars show approximate 95% confidence intervals (mean $\pm 1.96 \times$ standard error across patient queries) and percentage annotations indicate relative change for TrialGPT 2.0 compared with TrialGPT 1.0.

matches by the trial-specific recommendation system (Fig. 5c). The recommendation category assigned to each target trial is summarized in Supplementary Fig. 7.

Finally, on the public benchmarks (SIGIR and the TREC 2021 and TREC 2022 Clinical Trials tracks^{30–32}), which provide standardized relevance judgments for trial-ranking evaluation, TrialGPT 2.0 improved ranking performance over TrialGPT 1.0 across all four metrics (Fig. 5d). Averaged across the three benchmarks, TrialGPT 2.0 improved Precision@10 by 4.5%, normalized discounted cumulative gain at 10 (nDCG@10) by 5.7%, mean reciprocal rank (MRR) by 9.8% and mean average precision (MAP) by 10.2%. In addition, it reduced inference burden, increasing processing speed by a factor of 3.20 and reducing input and output token counts per patient-trial pair by 58% and 73%, respectively (Supplementary Note 4 and Supplementary Fig. 3b). Together, these results complement the real-world oncology evaluations by demonstrating improved recommendation-oriented performance, ranking quality, and computational efficiency under reproducible benchmark conditions.

Real-world deployment, runtime and usability of TrialGPT 2.0

Finally, we deployed TrialGPT 2.0 as a web-based interface for trial recommendation review across participating oncology and referral workflows. The interface accepts a patient clinical summary and uses the corresponding predefined local trial corpus and matching policy to generate ranked recommendations, displayed as trial cards grouped into Highly Recommended, Possible Match or Low Fit. Each trial card presents trial metadata and a structured rationale, including eligibility-supporting evidence, missing information and ineligibility concerns. Clinical reviewers and recruitment specialists can inspect recommendations by category, apply post-ranking filters and export matched results for review (Supplementary Fig. 4).

In runtime testing across 25 cases sampled from the five participating settings, TrialGPT 2.0 generated recommendations in 21.7 seconds on average, including retrieval when applicable, backend assessment, output parsing, result writing and score-sorted PDF generation. Runtime varied across settings, reflecting differences in case length, trial-corpus size and number of matched trials, but remained compatible with rapid generation of candidate trial lists for clinical review. In practical deployment, the interface supports recommendation generation from a nationwide corpus of more than 25,000 active clinical trials with sites in the United States. Additional runtime details are provided in Supplementary Note 5 and Supplementary Fig. 5.

Clinical reviewer acceptance was assessed using an 8-item questionnaire covering usability, workflow fit, perceived efficiency, clarity, decision confidence, usefulness, content safety, and overall satisfaction on a 5-point Likert scale. Overall, responses were favorable, with 85 of 104 item-level ratings scored as 4 or 5. Content safety received the highest mean rating of 4.92, with all responses favorable. Overall satisfaction and usability each received a mean score of 4.54, followed by clarity and workflow fit, with mean scores of 4.38 and 4.31, respectively. Ratings were more variable for perceived efficiency and decision confidence, with mean scores of 3.85 and 3.92, respectively, suggesting that perceived benefit may depend on workflow context, reviewer role, and continued user calibration. These findings support the operational feasibility of TrialGPT 2.0 across heterogeneous oncology and referral workflows, while highlighting the need for continued workflow integration and reviewer training during deployment (Supplementary Note 6, Supplementary Table 2 and Supplementary Fig. 6).

Discussion

In this multicenter study, we present TrialGPT 2.0, a real-world system for AI-assisted patient-trial recommendation, and evaluate it through retrospective review, prospective precision-oncology tumor-board assessment, reproducible benchmarking, and deployment. The clinical evaluation was predominantly oncology-focused, spanning disease-specific trial navigation, radiation-oncology treatment-fit assessment, institutional oncology referral screening, and molecularly guided tumor-board review. Across retrospective cohorts, TrialGPT 2.0 identified clinician-recommended trials within short ranked lists and reduced clinician screening time while preserving high agreement with consensus reference labels. In the prospective UIC POTB workflow, it expanded trial opportunities by identifying final recommendations for cases whose routine review yielded no final recommendation. On NIH-TrialBench and public benchmarks, it improved target-trial recovery and ranking performance relative to TrialGPT 1.0. TrialGPT 2.0 has been deployed through a secure web-based interface for trial review in real-world clinical workflows. The web-based interface generated recommendations within a short turnaround and received generally favorable ratings from clinical reviewers, supporting the operational feasibility of AI-assisted trial matching in routine trial-review workflows.

These findings address two gaps that have limited the clinical value of prior AI-assisted trial matching. First, as in much of medical AI, prior trial-matching methods have been evaluated largely using curated benchmarks that measure technical performance rather than real-world clinical contribution^{26,28}. By instead assessing TrialGPT 2.0 on real clinical data through clinician adjudication and prospective review in an active precision-oncology workflow, we show that AI-identified trials can contribute to real clinical decisions. To make such evaluation reproducible under the privacy constraints on real clinical notes, we developed NIH-TrialBench, which pairs clinician-authored synthetic vignettes built around known target trials with clinician-adjudicated labels and a realistic, workflow-defined trial search space.

Second, identifying eligible trials is not sufficient for clinical recommendation. Across the evaluations, reviewers explicitly separated trials that were Recommended from those that were Eligible but Not Recommended, reflecting differences in clinical appropriateness and priority that eligibility assessment alone does not capture. Accordingly, TrialGPT 2.0 produces graded recommendation categories and a fit score rather than a flat set of eligible trials, jointly assessing formal eligibility, patient-trial fit, and workflow-specific clinical priorities. Its fit score responds to clinically important mismatches in diagnosis or histology, disease status, biomarker or pathway match, line of therapy, prior treatment, and cohort or arm applicability, and is higher for clinician-recommended trials (Supplementary Fig. 1). Importantly, clinician adjudication showed that technically eligible trials may still be deprioritized because of treatment intent, limited expected benefit, preference for standard-of-care or more specific targeted options, treatment-sequence conflicts, feasibility, or patient-specific clinical context. Patient-trial matching is therefore better framed and evaluated as a recommendation task that extends beyond eligibility classification.

This study has several limitations. First, we evaluated early-stage trial-matching outcomes, including recommendation retrieval, clinician selection, agreement and screening time, but did not assess downstream recruitment outcomes such as patient contact, referral, formal eligibility screening, consent, enrollment, accrual or patient outcomes. These downstream processes require longer follow-up and are influenced by multiple factors beyond trial matching, including patient preferences, changes in clinical status, trial availability, logistical barriers and site-level recruitment procedures. Assessing the effect of TrialGPT 2.0 on actual enrollment and patient outcomes was

therefore beyond the scope of this study. In addition, the prospective evaluation was conducted in a single POTB workflow over a limited period with a modest number of cases; larger prospective evaluations across institutions, disease programs and recruitment workflows are needed to assess generalizability and downstream impact^{38,39}, ideally measuring trials discussed with patients, screening referrals, enrollment outcomes and accrual efficiency²⁸. Encouragingly, the original TrialGPT framework¹⁰ has already been independently deployed for real-world eligibility screening at an external institution⁴⁰, supporting the feasibility of adapting the approach outside its original development setting. Second, the evaluated trial corpora and clinical workflows were predominantly U.S.-based. International deployment would require integrating additional regional and international trial data sources, such as the EU Clinical Trials Information System⁴¹, the WHO International Clinical Trials Registry Platform⁴² and the Japan Registry of Clinical Trials⁴³. Because TrialGPT 2.0 is designed to be transferable across standardized candidate trial corpora, incorporating such sources is feasible and would broaden the system's applicability beyond the settings evaluated here. Third, the main evaluation used GPT-4.1 as the backbone model, which was available when system development and evaluation began but may be deprecated over time. Since TrialGPT 2.0 can be interpreted as a versioned AI-assisted system rather than a single static model, the backbone can be replaced with newer or alternative models. Supplementary analyses show that it transfers to GPT-5.4 and Gemini-3.1-pro-preview with a similar overall pattern of performance (Supplementary Fig. 3). Fourth, error analysis revealed failure modes in both directions. TrialGPT 2.0 occasionally applied eligibility criteria too literally, missing trials that clinicians retained as reasonable options by drawing on broader oncologic context, such as the clinical interpretation of prior treatment or trial feasibility (Fig. 4). Conversely, among trials that TrialGPT 2.0 rated Highly Recommended but clinicians did not recommend, discordances reflected condition, biomarker or protocol mismatches, as well as clinical-priority concerns such as limited expected benefit or treatment-sequence conflict (Supplementary Fig. 2). Reducing these residual errors, particularly in complex multi-arm trials and cases with incomplete or clinically ambiguous information, is an important direction for future work.

Our findings show that AI-assisted matching can already identify additional trial options that clinicians choose to act on. Realizing its full value, however, will require integrating these evolving systems into rigorous clinical workflows and pairing them with large-scale, outcome-based evaluation. With ongoing performance monitoring and studies that measure patient outcomes, AI-assisted matching could progress from suggesting candidate trials to producing real improvements in cancer trial access, helping more patients reach the trials for which they are eligible, earlier and more efficiently.

Methods

TrialGPT 2.0 system design

TrialGPT 2.0 takes three inputs for each matching task: patient context, a predefined local trial corpus and a local matching policy. Patient context can be provided as raw clinical records or as a clinician-prepared narrative summary. The local trial corpus comprises the predefined set of candidate trials to be considered for matching in a given recruitment setting, including trial metadata, eligibility criteria and cohort or arm information. The local matching policy specifies clinical preferences that guide matching and ranking, including clinical relevance weighting, disease- or biomarker-specific priorities and therapeutic trial preference. These preferences guide prioritization but do not override explicit eligibility criteria.

TrialGPT 2.0 uses a pluggable retrieval component to support real-time matching across trial corpora of different sizes. For large trial pools, detailed backend assessment of every trial can be computationally inefficient and unsuitable for real-time review. In these settings, the default retriever follows the hybrid-fusion retrieval strategy of the original TrialGPT¹⁰ framework. Briefly, the retriever generates patient-derived search queries from the patient context, applies both lexical and semantic retrieval to the local trial corpus, and merges query-level results using reciprocal rank fusion to produce a candidate shortlist. In our implementation, lexical retrieval is based on BM25⁴⁴ and semantic retrieval is based on MedCPT⁴⁵, consistent with the original TrialGPT retrieval design. For smaller trial pools, defined in this study as fewer than 1,500 candidate trials, TrialGPT 2.0 passes the full predefined trial corpus directly to the backend. The retrieval component determines which trials proceed to detailed assessment, whereas the backend assigns the final fit score, confidence estimate, recommendation category and rank.

The TrialGPT 2.0 backend generates a single trial-level assessment for each candidate trial. When a trial includes multiple cohorts or arms, the backend evaluates whether the patient fits any clinically meaningful option and produces one overall recommendation for the trial. The assessment jointly considers formal eligibility, clinical relevance and uncertainty in the available patient information. Clear violations of required eligibility criteria lower the fit score and appear as ineligible reasons, whereas missing information that could typically be obtained during screening lowers confidence rather than automatically indicating ineligibility. When line-of-therapy information is available, the backend uses it to determine whether the patient's prior treatment history matches the treatment setting required by the trial or relevant cohort.

For each candidate trial, TrialGPT 2.0 returns a strict structured output containing eligible reasons, missing information, ineligible reasons, rationale, fit score and confidence estimate. Eligible reasons summarize criteria that the patient appears to meet; ineligible reasons summarize explicit mismatches or exclusion concerns; missing information captures key unknowns that affect assessment; and the rationale provides an overall explanation of the recommendation. The fit score is an integer from 0 to 100 and represents the overall strength of patient-trial alignment. Recommendation categories are assigned from the fit score: scores greater than 90 are classified as Highly Recommended, scores from 80 to 90 as Possible Match and scores below 80 as Low Fit. The confidence estimate ranges from 0 to 1 and reflects the reliability of the assessment, accounting for missing or uncertain information. The final rank orders trials primarily by fit score; when trials have the same fit score, confidence serves as a tie-breaking criterion, with higher-confidence assessments ranked above lower-confidence assessments.

Unless otherwise specified, all main analyses used GPT-4.1⁴⁶, accessed through the Azure OpenAI

Service, as the backbone model. Model decoding settings, code version, prompt and backbone model were fixed across evaluations. Analyses with alternative backbone models (GPT-5.4⁴⁷ and Gemini-3.1-pro-preview⁴⁸) are reported in Supplementary Fig. 3.

Matching policies were developed with clinician input to reflect the review priorities and available patient information in each recruitment or evaluation setting. Consequently, the trial-fit factors represented in the policies, and the depth with which they were assessed, varied across workflows (Supplementary Note 1 and Supplementary Table 1). For the NIH-TrialBench evaluation, TrialGPT 2.0 uses a general matching policy that is fixed before evaluation and is not refined using NIH-TrialBench vignettes or annotations. For retrospective cohorts, local matching policies are refined using clinician feedback from cases other than the case being evaluated. For each retrospective evaluation case, the policy is fixed before inference, and any clinician feedback or error analysis from that same case is not used to revise the policy that generated its reported result.

Because LLM backbones, trial registries and local trial portfolios can change over time, we evaluated TrialGPT 2.0 as a versioned AI-assisted system. For NIH-TrialBench, retrospective clinical-note cohorts, public benchmark analyses and runtime experiments, the code version, prompt version, backbone model, decoding settings, retrieval configuration, trial-corpus or benchmark candidate set and matching policy were fixed before inference. For the prospective POTB evaluation, the system used the current weekly refreshed UIC trial-corpus snapshot available at the time of each case review, while the code version, prompt version, backbone model, decoding settings, retrieval configuration and matching policy were recorded. Reported results therefore apply to the TrialGPT 2.0 configurations and trial-corpus snapshots used for each evaluation. Substantive changes to the backbone model, prompts, retrieval module, trial corpus or matching policy should be treated as a new system version and undergo local revalidation before clinical use.

Study cohorts

The study included three cohort types: de-identified retrospective clinical notes, prospective Precision Oncology Tumor Board (POTB) cases and clinician-authored synthetic vignettes. These cohort types supported complementary evaluation of TrialGPT 2.0 across real clinical documentation, an active clinical trial-review workflow and reproducible benchmark conditions.

The retrospective evaluation used 288 de-identified cases from five recruitment pathways: patient call intake notes from the NIH Office of Patient Recruitment (OPR), oncology patient advocacy intake notes from the Cholangiocarcinoma Foundation (CCF), tumor-board referrals from the University of Illinois Cancer Center (UIC), oncology consultation notes from the University of Pittsburgh Medical Center (UPMC) and oncology referral notes from the National Cancer Institute (NCI). For each retrospective cohort, the contributing site or program defined a candidate trial corpus before TrialGPT 2.0 inference. These corpora reflected the trials considered relevant to that recruitment setting and were fixed during retrospective inference and clinician review. Trial search-space size varied across cohorts, from 9 to 1,871 candidate trials.

The prospective evaluation was conducted in the UIC Precision Oncology Tumor Board (POTB) workflow. Prospective cases entered the workflow through routine tumor-board referral and were reviewed using de-identified case context, including clinical history, molecular findings, prior therapies, and the clinical question. Unlike the fixed retrospective trial-corpus snapshots, the prospective workflow used the current UIC trial-corpus snapshot available at the time of each case review, reflecting weekly updates to trial availability and recruitment status.

To support reproducible evaluation, we created NIH-TrialBench, a clinician-authored synthetic benchmark based on actual NIH trials. NIH-TrialBench was developed from target-trial-informed vignettes contributed by 24 investigators, with target trials spanning 11 NIH Institutes and Centers. The main NIH-TrialBench evaluation set included 126 synthetic vignettes and a predefined corpus of 1,373 candidate trials from the NIH Clinical Center’s *Search the Studies* registry³⁵. Each vignette corresponded to one target trial, defined as the intended match for the synthetic patient. Because the vignettes are synthetic, NIH-TrialBench supports reproducible benchmarking without exposing private patient information.

Cohort-specific review strategies are summarized in Fig. 2a. Top-ranked review was used for NIH-TrialBench and the OPR, CCF, UIC and UPMC retrospective cohorts. Exhaustive review was used for the NCI retrospective cohort because each case had only nine predefined candidate trials, making complete review feasible. Prospective review was used in the UIC POTB workflow, where TrialGPT 2.0-generated candidate trials were reviewed during ongoing trial recommendation.

We manually reviewed clinician annotations for quality control before analysis, including checks for completeness and consistency between annotation comments and assigned labels. Annotation issues identified during quality control were resolved before metric calculation.

We characterized disease-area coverage using first-level disease categories from Medical Subject Headings (MeSH)³³. For each retrospective clinical note, prospective POTB case and NIH-TrialBench vignette, GPT-5.4⁴⁷, accessed through the Azure OpenAI Service, mapped the patient context to up to three predefined MeSH disease categories, ordered by relevance. Model outputs were restricted to the predefined category list and manually reviewed for quality control. We summarized disease-area coverage as the number of cases in each cohort assigned to each MeSH disease category (Fig. 2b). Age group and sex were obtained from case metadata or vignette text when available; sex was recorded as unknown when not specified (Fig. 2c).

Retrospective evaluation

Retrospective clinical-note evaluation assessed whether TrialGPT 2.0 recommendations aligned with clinician judgment and whether TrialGPT 2.0 assistance reduced clinician screening time. The retrospective analysis included two review strategies. Top-ranked review was applied to the OPR, CCF, UIC, and UPMC cohorts and evaluated whether highly ranked TrialGPT 2.0 recommendations were judged eligible or recommended by clinicians. Exhaustive review was applied to the NCI cohort, where each case had a predefined search space of nine candidate trials, enabling complete-pool agreement analysis and a counterbalanced clinician-assistance experiment.

For top-ranked review, TrialGPT 2.0 ranked the predefined candidate trial corpus for each case. Clinicians reviewed up to ten highest-ranked trials categorized by TrialGPT 2.0 as Highly Recommended and assigned each reviewed patient-trial pair one of three labels: Recommended, Eligible but Not Recommended, or Ineligible. Recommended indicates that the patient appears eligible and that the trial is clinically advisable for consideration. Eligible but Not Recommended indicates that the patient appears technically eligible but that the trial is not clinically advisable or not a priority in the current clinical context. Ineligible indicates that the patient does not meet trial eligibility criteria or has a clear incompatibility with the trial. We evaluated top-ranked review performance at $K \in \{1, 3, 5, 10\}$. For these metrics, 95% confidence intervals were computed by percentile bootstrap over cases with 10,000 resamples (Fig. 3b).

Let q denote a case and let $\mathcal{Q}_{\text{top}}$ denote the set of retrospective cases in the top-ranked review analysis. Let $S_q(K)$ denote the first $\min(K, n_q)$ annotated trials for case q , where n_q is the number of annotated Highly Recommended trials for that case. Let $y_{q,t}$ denote the clinician label for trial t in case q , with

$$y_{q,t} \in \{\text{Recommended, Eligible but Not Recommended, Ineligible}\}. \quad (1)$$

We define the set of clinically eligible labels as

$$\mathcal{E} = \{\text{Recommended, Eligible but Not Recommended}\}. \quad (2)$$

Hit rate@ K measures case-level success in identifying at least one clinician-recommended trial within the evaluated top- K list:

$$\text{Hit rate@}K = \frac{1}{|\mathcal{Q}_{\text{top}}|} \sum_{q \in \mathcal{Q}_{\text{top}}} \mathbb{I}(\exists t \in S_q(K) : y_{q,t} = \text{Recommended}). \quad (3)$$

Because clinician review was restricted to model-selected Highly Recommended trials, hit rate@ K measures whether the evaluated top- K list contained at least one clinician-recommended trial and should not be interpreted as recall over the full candidate trial corpus.

$$\text{Recommended precision@}K = \frac{1}{|\mathcal{Q}_K|} \sum_{q \in \mathcal{Q}_K} \frac{1}{|S_q(K)|} \sum_{t \in S_q(K)} \mathbb{I}(y_{q,t} = \text{Recommended}). \quad (4)$$

Eligible precision@ K measures the fraction of evaluated top- K recommendations judged clinically eligible by clinicians:

$$\text{Eligible precision@}K = \frac{1}{|\mathcal{Q}_K|} \sum_{q \in \mathcal{Q}_K} \frac{1}{|S_q(K)|} \sum_{t \in S_q(K)} \mathbb{I}(y_{q,t} \in \mathcal{E}), \quad (5)$$

where $\mathcal{Q}_K = \{q : |S_q(K)| > 0\}$.

The NCI cohort represents a distinct evaluation setting because clinicians exhaustively reviewed every trial in a small, curated portfolio of nine potentially relevant oncology trials. Since the search space had already been narrowed to trials relevant to the workflow, the principal task was eligibility triage. *Eligible* indicated that no current eligibility barrier had been identified and that the trial could proceed to further consideration or screening. *Subject-to-change ineligible* indicated that eligibility could change with additional information or modification of a clinical factor, whereas *Definitely ineligible* indicated a clear, non-resolvable incompatibility. For the three-class agreement analysis, we mapped *Eligible* to TrialGPT 2.0’s Highly Recommended category, *Subject-to-change ineligible* to Possible Match and *Definitely ineligible* to Low Fit. In this focused setting, mapping *Eligible* to Highly Recommended is not contradictory because Eligible functioned as the positive action category for proceeding to further consideration. This mapping was used only to align the original clinician labels with the system’s three output categories. We denote this label set as

$$\mathcal{C} = \{\text{Highly Recommended, Possible Match, Low Fit}\}. \quad (6)$$

Each NCI patient-trial pair received three initial human annotations. Before the assistance experiment, an independent clinician blinded to TrialGPT 2.0 outputs annotated all 900 patient-trial

pairs. The same pairs were then reviewed by two additional clinicians in a counterbalanced TrialGPT 2.0 assistance experiment. The 100 cases were divided into two groups of 50: one clinician reviewed the first group with TrialGPT 2.0 assistance and the second group without assistance, while the other clinician reviewed the groups under the opposite assistance assignments. Thus, each patient-trial pair received one independent blinded annotation, one annotation with TrialGPT 2.0 assistance, and one annotation without assistance. When all three initial annotations agreed, their shared label was used as the consensus reference label. When the annotations disagreed, a fourth clinician reviewed the discordant pair and made the final determination.

The labels resulting from this process were used as consensus reference labels for agreement analysis. Let $r_{q,t}$ denote the consensus reference label for trial t in case q . Let $m_{q,t}$ denote the TrialGPT 2.0 recommendation category, $h_{q,t}^+$ denote the clinician label assigned with TrialGPT 2.0 assistance, and $h_{q,t}^-$ denote the clinician label assigned without TrialGPT 2.0 assistance, with

$$r_{q,t}, m_{q,t}, h_{q,t}^+, h_{q,t}^- \in \mathcal{C}. \quad (7)$$

We compared three assigned-label sources with the consensus reference labels: TrialGPT 2.0 alone, clinician review with TrialGPT 2.0 assistance, and clinician review without TrialGPT 2.0 assistance. For each source s , let $z_{q,t}^{(s)}$ denote its assigned label, where

$$z_{q,t}^{(s)} \in \{m_{q,t}, h_{q,t}^+, h_{q,t}^-\}. \quad (8)$$

We computed a three-class confusion matrix for each source. Rows correspond to assigned labels and columns correspond to consensus reference labels:

$$C_{a,b}^{(s)} = \sum_{q \in \mathcal{Q}_{\text{NCI}}} \sum_{t \in \mathcal{T}_q} \mathbb{I}(z_{q,t}^{(s)} = a, r_{q,t} = b), \quad (9)$$

where $a, b \in \mathcal{C}$, $\mathcal{Q}_{\text{NCI}}$ denotes the NCI retrospective cohort and $\mathcal{T}_q$ denotes the nine candidate trials for case q .

Exact agreement for source s is defined as the proportion of patient-trial pairs for which the assigned label matches the consensus reference label:

$$\text{Agreement}^{(s)} = \frac{\sum_{c \in \mathcal{C}} C_{c,c}^{(s)}}{\sum_{a \in \mathcal{C}} \sum_{b \in \mathcal{C}} C_{a,b}^{(s)}}. \quad (10)$$

Class-specific recall for reference category c is defined as the proportion of consensus-reference pairs in category c that the source assigns to the same category:

$$\text{Recall}_c^{(s)} = \frac{C_{c,c}^{(s)}}{\sum_{a \in \mathcal{C}} C_{a,c}^{(s)}}. \quad (11)$$

We also summarized severe category reversals between Highly Recommended and Low Fit. Severe over-recommendation was defined as a consensus-reference Low Fit pair assigned as Highly Recommended:

$$\text{OverRec}^{(s)} = C_{\text{Highly Recommended}, \text{Low Fit}}^{(s)}. \quad (12)$$

Severe under-recommendation was defined as a consensus-reference Highly Recommended pair assigned as Low Fit:

$$\text{UnderRec}^{(s)} = C_{\text{Low Fit}, \text{Highly Recommended}}^{(s)}. \quad (13)$$

The total number of severe high-low reversals was calculated as

$$\text{Reversal}^{(s)} = \text{OverRec}^{(s)} + \text{UnderRec}^{(s)}. \quad (14)$$

We analyzed clinician screening time in the NCI counterbalanced TrialGPT 2.0 assistance experiment. Screening time was recorded at the patient-trial-pair level and represented clinician decision time for each candidate trial. Patient-context review time and automated TrialGPT 2.0 PDF report-generation time were not included. For each case q , total screening time with TrialGPT 2.0 assistance was calculated as

$$T_q^+ = \sum_{t \in \mathcal{T}_q} \tau_{q,t}^+, \quad (15)$$

where $\tau_{q,t}^+$ denotes the trial-level screening time recorded under TrialGPT 2.0 assistance. Total screening time without TrialGPT 2.0 assistance was calculated as

$$T_q^- = \sum_{t \in \mathcal{T}_q} \tau_{q,t}^-, \quad (16)$$

where $\tau_{q,t}^-$ denotes the corresponding trial-level screening time without TrialGPT 2.0 assistance.

Mean total screening time was summarized across cases for the assisted and unassisted conditions. Percentage reduction in screening time was calculated as

$$\text{Reduction} = \frac{\overline{T^-} - \overline{T^+}}{\overline{T^-}} \times 100\%. \quad (17)$$

Timing results were summarized overall and within each 50-case counterbalanced group. Error bars in Fig. 3d show 95% confidence intervals of the mean, computed by percentile bootstrap over cases with 10,000 resamples. Statistical comparisons between assisted and unassisted screening times used two-sided Wilcoxon signed-rank tests at the case level. Total screening time was also compared between the two clinicians using a two-sided Mann-Whitney U test to assess whether the reduction could be explained by systematic clinician-level speed differences.

Prospective evaluation

We prospectively evaluated TrialGPT 2.0 in the University of Illinois Cancer Center Precision Oncology Tumor Board (POTB) trial-matching workflow. The analysis included POTB-reviewed cases from February through July 2026 for which TrialGPT 2.0 outputs, routine-workflow-identified trials and final POTB clinical trial recommendations were available. One referred case from April was excluded because the patient was deceased before the POTB session and the case was not included in the final April case list.

For each referred case, the POTB materials summarized the clinical context, molecular findings, prior therapies, and the POTB review question. Trial identification proceeded in parallel: clinical reviewers identified candidate trials through the routine POTB workflow, while TrialGPT 2.0 generated candidate trial recommendations using the UIC trial corpus and matching policy. The two candidate-trial lists were subsequently reviewed together, and clinicians determined the final POTB clinical trial recommendation list. Source attribution reflected whether each final recommendation had been identified in the independently generated TrialGPT 2.0 list, the routine-workflow list or both.

The primary prospective endpoint quantified TrialGPT 2.0’s contribution to final clinician-selected trial recommendations at both the case and recommendation levels. At the case level, each reviewed case was assigned to one of three mutually exclusive categories: no final trial recommendation, final trial options expanded by TrialGPT 2.0 or final trial options unchanged by TrialGPT 2.0. A case was classified as expanded when the routine workflow identified no final trial recommendation but TrialGPT 2.0 contributed at least one trial that clinicians selected for the final recommendation list. A case was classified as unchanged at this case-level opportunity threshold when the routine workflow had already identified at least one final recommendation, regardless of whether TrialGPT 2.0 contributed additional recommendations. Cases for which neither source contributed a final recommendation were classified as having no final trial recommendation. We also reported, among cases with at least one final recommendation, the proportion for which TrialGPT 2.0 contributed at least one recommendation and the proportion for which TrialGPT 2.0 was the sole source of the final recommendations (Fig 4a).

At the recommendation level, each final recommendation was attributed to one of three mutually exclusive sources: TrialGPT 2.0 only, both TrialGPT 2.0 and the routine workflow, or the routine workflow only. TrialGPT 2.0-only recommendations were generated by TrialGPT 2.0 and selected by clinicians for the final recommendation list but were not identified through the routine workflow. Recommendations identified independently by both sources were classified as overlapping recommendations. Routine-workflow-only recommendations were identified through the routine workflow but not by TrialGPT 2.0. TrialGPT 2.0-associated recommendations comprised those attributed to TrialGPT 2.0 only or to both sources. Counts and proportions were summarized overall and by POTB month (Fig 4b).

The secondary prospective endpoint characterized all TrialGPT 2.0-generated candidate trials using UIC clinician labels. Each candidate trial was labeled as Recommended, Eligible but Not Recommended or Ineligible. Recommended was assigned when the candidate trial was selected for the final POTB recommendation list. Eligible but Not Recommended was assigned when the trial was considered potentially eligible but was not prioritized for final recommendation in the prospective workflow. Ineligible was assigned when the trial was considered incompatible with the patient context or trial criteria. We reported the number and proportion of TrialGPT 2.0-generated candidate trials in each category overall and by POTB month (Fig 4c).

Reproducible benchmarking

Reproducible benchmarking evaluation included NIH-TrialBench and previously released public patient-trial matching benchmarks. NIH-TrialBench was used to assess clinician-adjudicated recommendation quality and intended target-trial recovery under synthetic benchmark conditions. Public benchmarks were used to assess standardized eligibility-oriented ranking performance and computational efficiency in comparison with the original TrialGPT framework.

For NIH-TrialBench top-ranked review, TrialGPT 2.0 ranked the predefined NIH-TrialBench candidate trial corpus for each vignette. Clinicians reviewed up to ten highest-ranked trials categorized by TrialGPT 2.0 as Highly Recommended and assigned each reviewed patient-trial pair one of three labels: Recommended, Eligible but Not Recommended or Ineligible. This analysis used the same top-ranked review metric definitions as the retrospective top-ranked review, including eligible precision@ K , recommended precision@ K and hit rate@ K for $K \in \{1, 3, 5, 10\}$, with the case set restricted to NIH-TrialBench vignettes. For these metrics, 95% confidence intervals were computed by percentile bootstrap over vignettes with 10,000 resamples (Fig. 5b).

NIH-TrialBench also enabled target-trial assessment because each vignette has a prespecified target trial representing the intended match for the synthetic patient. After TrialGPT 2.0 scored the predefined NIH-TrialBench trial corpus, we located the target trial in the system output and recorded its assigned recommendation category. We summarized the number and proportion of target trials categorized as Highly Recommended, Possible Match or Low Fit.

For the target-trial Recall@10 analysis, each evaluated method ranked candidate trials for each NIH-TrialBench vignette. For TrialGPT 2.0, ranks were generated from the predefined NIH-TrialBench candidate trial corpus. For the off-the-shelf LLM baselines, we queried the ChatGPT web interface using GPT-5.5 Thinking³⁶ and the Gemini web interface using Gemini-3.5 Flash with Extended Thinking³⁷. Baseline queries were conducted on May 22–23, 2026. For each vignette, we started a new temporary chat session to minimize the effect of prior conversation history and submitted a standardized instruction asking the model to search the NIH Clinical Center *Search the Studies* registry³⁵ and return only the NCT identifiers of the ten recommended trials for the patient vignette. We used Extended Thinking for Gemini-3.5 Flash because the Standard thinking setting did not reliably return valid trial recommendations under this instruction. Returned NCT identifiers were parsed in the order provided by the model. A target trial was counted as recovered if its NCT identifier appeared among the ten returned identifiers (Supplementary Fig. 7).

Target-trial Recall@10 was calculated separately for each contributing NIH Institute or Center and summarized as an unweighted macro-average across Institutes and Centers. Let $\mathcal{G}$ denote the set of contributing NIH Institutes and Centers, let $Q_g \subseteq Q_{\text{syn}}$ denote the set of NIH-TrialBench vignettes contributed by institute or center g , and let m denote an evaluated method.

Institute- or center-specific target-trial Recall@10 was defined as

$$\text{Recall@10}_{m,g} = \frac{1}{|Q_g|} \sum_{q \in Q_g} \mathbb{I}[\text{rank}_{m,q}(t_q) \leq 10], \quad (18)$$

and macro-averaged target-trial Recall@10 was defined as

$$\text{Recall@10}_{m,\text{macro}} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \text{Recall@10}_{m,g}. \quad (19)$$

Here, t_q denotes the prespecified target trial for vignette q , and $\text{rank}_{m,q}(t_q)$ denotes the position of that target trial in the ranked list produced by method m . When the target trial was not returned among the ten recommendations, its rank was treated as greater than 10. The Avg bars in Fig. 5c report $\text{Recall@10}_{m,\text{macro}}$, giving each contributing NIH Institute or Center equal weight regardless of the number of vignettes it contributed.

For public benchmark evaluation, we evaluated TrialGPT 2.0 on the 2016 Special Interest Group on Information Retrieval (SIGIR) patient-trial matching collection and the TREC 2021 and TREC 2022 Clinical Trials tracks^{30–32}. These benchmarks provide standardized relevance judgments for patient-trial matching.

The public benchmark analysis evaluated trial-level matching and ranking, not first-stage retrieval. The TrialGPT 2.0 retrieval module was not applied in this stage. For each patient, the trial search space consisted of all trials with available relevance judgments for that patient. TrialGPT 1.0 and TrialGPT 2.0 each scored every labeled patient-trial pair in this search space, and the scored trials were sorted to produce a ranked list for metric calculation. This design avoids treating unlabeled trials as negative examples and isolates the comparison to trial-level assessment and ranking. The evaluated candidate pools included 58 cases and 3,835 judged patient-trial pairs for SIGIR, 75

cases and 35,832 judged pairs for TREC 2021, and 50 cases and 35,394 judged pairs for TREC 2022.

Benchmark relevance judgments were represented as three-level graded scores. In SIGIR, scores of 0, 1 and 2 correspond to irrelevant, potential and eligible trials, respectively. For TREC 2021 and TREC 2022, we used the official graded relevance judgments, in which 0 denotes non-relevant, 1 denotes excluded and 2 denotes eligible. For Precision@10 and nDCG@10, graded relevance scores were used directly. For mean reciprocal rank (MRR) and mean average precision (MAP), graded relevance was converted to binary relevance by treating eligible trials, corresponding to score 2, as relevant. Patient-trial pairs without relevance judgments were excluded from metric calculation.

Let q denote a patient, i denote the rank position of a trial in the model-generated ranked list, $r_{q,i}$ denote the graded relevance score of the trial ranked at position i for patient q , and $r_{q,i}^*$ denote the graded relevance score at position i in the ideal ranking for that patient. The maximum relevance score is $r_{\max} = 2$.

Precision@10 is computed as the normalized graded relevance of the top ten ranked trials:

$$\text{Precision@10}(q) = \frac{\sum_{i=1}^{10} r_{q,i}}{10 \times r_{\max}}. \quad (20)$$

nDCG@10 is computed as

$$\text{nDCG@10}(q) = \frac{\text{DCG@10}(q)}{\text{IDCG@10}(q)}, \quad (21)$$

where

$$\text{DCG@10}(q) = \sum_{i=1}^{10} \frac{r_{q,i}}{\log_2(i+1)} \quad (22)$$

and

$$\text{IDCG@10}(q) = \sum_{i=1}^{10} \frac{r_{q,i}^*}{\log_2(i+1)}. \quad (23)$$

We set nDCG@10 to 0 for cases with IDCG@10(q) = 0.

For MRR and MAP, eligible trials with $r_{q,i} = 2$ were treated as relevant. Reciprocal rank for patient q is defined as

$$\text{RR}(q) = \frac{1}{\min\{i : r_{q,i} = 2\}}, \quad (24)$$

with reciprocal rank set to 0 if no eligible trial appeared in the ranked list. MRR is the mean of $\text{RR}(q)$ across cases.

Average precision for patient q is computed as

$$\text{AP}(q) = \frac{1}{M_q} \sum_{i=1}^{N_q} \text{P@}i(q) \cdot \mathbb{I}(r_{q,i} = 2), \quad (25)$$

where N_q is the number of ranked trials for patient q , $M_q = \sum_{i=1}^{N_q} \mathbb{I}(r_{q,i} = 2)$, $\text{P@}i(q)$ is binary precision among the top i ranked trials using the same $r = 2$ relevance threshold, and $\mathbb{I}(r_{q,i} = 2)$

indicates whether the trial at rank i is eligible. We set average precision to 0 for cases with $M_q = 0$. MAP is the mean of $AP(q)$ across cases.

Dataset-level metrics were averaged across cases. The overall benchmark average was calculated as the arithmetic mean of the three dataset-level values from SIGIR, TREC 2021 and TREC 2022. For individual benchmark results, approximate 95% confidence intervals were computed as the mean ± 1.96 standard errors across patient queries. For average bars across benchmarks, error bars summarize variation across the three dataset-level means. Relative change for TrialGPT 2.0 compared with TrialGPT 1.0 was calculated as

$$\text{Relative change} = \frac{\text{Metric}_{\text{TrialGPT 2.0}} - \text{Metric}_{\text{TrialGPT 1.0}}}{\text{Metric}_{\text{TrialGPT 1.0}}} \times 100\%. \quad (26)$$

Computational efficiency was measured at the patient-trial-pair level using the same judged candidate pools evaluated for ranking. For each scored pair, we recorded latency and model token usage. For TrialGPT 1.0, latency and token counts were summed across the criterion-level matching and aggregation calls for the pair. For TrialGPT 2.0, latency and token counts were recorded from the corresponding trial-level scoring call. Processing speed was defined as the inverse of the mean per-pair latency and reported as patient-trial pairs scored per second. Input tokens included all tokens submitted to the backbone model, and output tokens included all tokens generated by the model. Average speed, input tokens and output tokens were calculated across patient-trial pairs within each benchmark.

The main public benchmark comparison used GPT-4.1⁴⁶ as the backbone model for both TrialGPT 1.0 and TrialGPT 2.0. Additional analyses with alternative backbone models are reported in Supplementary Fig. 3.

Real-world deployment evaluation

To facilitate structured clinician review, the TrialGPT 2.0 interface accepts patient context, initiates matching against the predefined local trial corpus and displays the resulting matched-trial list. The interface uses the same TrialGPT 2.0 backend, matching policy and trial-corpus versions as the corresponding evaluation setting. Interface-level filtering, sorting and PDF export are applied after TrialGPT 2.0 generates the matched-trial output and do not change the model-generated recommendation category, rank or structured explanation used for evaluation.

In participating deployment settings, local trial search spaces are refreshed weekly to reflect changes in trial availability, recruitment status, cohort or arm availability and site-specific trial lists. Each recommendation run uses the local trial-corpus snapshot available at the time of matching. For NIH-TrialBench, retrospective clinical-note cohorts, public benchmark analyses and runtime experiments, the trial-corpus or benchmark candidate set was fixed before inference and recorded, so later weekly updates did not alter the reported results. For the prospective POTB evaluation, TrialGPT 2.0 used the current UIC trial-corpus snapshot available at the time of each case review, reflecting the weekly update process used in the live workflow; the snapshot used for each prospective run was recorded for auditability.

We measured recommendation-generation time in the interface deployment environment. Timing started when the user initiated matching after entering the patient context and ended when the matched-trial list was displayed in the interface. The measurement included retrieval when applicable, trial loading, backend LLM-based assessment, score aggregation and ranking, output parsing, result writing and score-sorted PDF generation. It excluded patient-context entry time,

manual clinician review and any post-generation human actions. For the runtime experiment, we sampled five cases from each evaluation source and ran each case three times using the corresponding site-specific trial corpus and matching configuration, in the interface runtime environment using 128 worker processes. The three repeated runs were averaged to obtain a case-level runtime estimate, and runtime was summarized across cases within each source as the mean and standard deviation (Supplementary Fig. 5).

Clinical reviewer acceptance was assessed using an 8-item questionnaire after reviewers used or reviewed TrialGPT 2.0 outputs in clinical trial matching workflows. The questionnaire assessed usability, workflow fit, perceived efficiency, clarity, decision confidence, usefulness, content safety and overall satisfaction using a 5-point Likert scale, where 1 indicated strongly disagree and 5 indicated strongly agree. We summarized item-level response distributions and favorable responses, defined as ratings of 4 or 5 (Supplementary Fig. 6).

Study approval and ethics oversight

The study was conducted in accordance with applicable institutional policies and regulatory requirements. Retrospective and prospective clinical-note cohorts were analyzed using de-identified, coded or anonymized data under applicable institutional approvals, waivers, data-use agreements or non-human-subjects research determinations at the contributing sites. TrialGPT 2.0 outputs were used to support clinician review; clinicians retained responsibility for final trial recommendations, referral, screening and enrollment decisions. Results are reported only in aggregate, and no directly identifiable participant information is included in the manuscript.

The National Cancer Institute retrospective cohort, NIH Office of Patient Recruitment cohort, and Cholangiocarcinoma Foundation cohort were analyzed using only de-identified or anonymized data and were determined to constitute Not Human Subjects Research (NHSR) under applicable federal regulations⁴⁹. Accordingly, informed consent was not required. No biospecimens were accessed, no participants were contacted for research purposes, and no additional clinical data were collected beyond information already available from routine referral, recruitment, or trial-matching workflows.

The University of Pittsburgh Medical Center cohort was analyzed under approval of the UPMC Institutional Review Board (STUDY20070273), with waivers of informed consent as applicable.

The University of Illinois Cancer Center retrospective and prospective Precision Oncology Tumor Board evaluations were conducted under UIC protocol STUDY2025-0124, with a waiver of documentation of informed consent and a waiver or alteration of HIPAA authorization.

NIH-TrialBench used clinician-authored synthetic vignettes constructed from target trials and did not use private patient records.

References

1. DiMasi, J. A., Grabowski, H. G. & Hansen, R. W. Innovation in the pharmaceutical industry: new estimates of r&d costs. *J. health economics* **47**, 20–33 (2016).
2. Sertkaya, A., Beleche, T., Jessup, A. & Sommers, B. D. Costs of drug development and research and development intensity in the us, 2000-2018. *JAMA network open* **7**, e2415445 (2024).

3. Unger, J. M., Vaidya, R., Hershman, D. L., Minasian, L. M. & Fleury, M. E. Systematic review and meta-analysis of the magnitude of structural, clinical, and physician and patient barriers to cancer clinical trial participation. *JNCI: J. Natl. Cancer Inst.* **111**, 245–255 (2019).
4. Penberthy, L. T., Dahman, B. A., Petkov, V. I. & DeShazo, J. P. Effort required in eligibility screening for clinical trials. *J. Oncol. Pract.* **8**, 365–370 (2012).
5. Speich, B. *et al.* Nonregistration, discontinuation, and nonpublication of randomized trials: A systematic review. *JAMA Netw. Open* **8**, e2524440 (2025).
6. Hauck, C. L., Kelechi, T. J., Cartmell, K. B. & Mueller, M. Trial-level factors affecting accrual and completion of oncology clinical trials: a systematic review. *Contemp. Clin. Trials Commun.* **24**, 100843 (2021).
7. Zhang, E. & DuBois, S. G. Early termination of oncology clinical trials in the united states. *Cancer Medicine* **12**, 5517–5525 (2023).
8. Schroen, A. T. *et al.* Achieving sufficient accrual to address the primary endpoint in phase iii clinical trials from us cooperative oncology groups. *Clin. Cancer Res.* **18**, 256–262 (2012).
9. Peterson, J. S. *et al.* Growth in eligibility criteria content and failure to accrue among national cancer institute (nci)-affiliated clinical trials. *Cancer Medicine* **12**, 4715–4724 (2023).
10. Jin, Q. *et al.* Matching patients to clinical trials with large language models. *Nat. communications* **15**, 9074 (2024).
11. Rybinski, M., Kusa, W., Karimi, S. & Hanbury, A. Learning to match patients to clinical trials using large language models. *J. Biomed. Informatics* **159**, 104734 (2024).
12. Chan, J. *et al.* Recommending clinical trials for online patient cases using artificial intelligence. *ArXiv arXiv-2504* (2025).
13. Abdallah, M. *et al.* Trialmatchai: an end-to-end ai-powered clinical trial recommendation system to streamline patient-to-trial matching. *Nat. Commun.* (2026).
14. Park, J. *et al.* Criteria2query 3.0: Leveraging generative large language models for clinical trial eligibility query generation. *J. biomedical informatics* **154**, 104649 (2024).
15. Klein, H. *et al.* Matchminer: an open-source platform for cancer precision medicine. *NPJ precision oncology* **6**, 69 (2022).
16. Zhang, X., Xiao, C., Glass, L. M. & Sun, J. Deepenroll: patient-trial matching with deep embedding and entailment prediction. In *Proceedings of the web conference 2020*, 1029–1037 (2020).
17. Beck, J. T. *et al.* Artificial intelligence tool for optimizing eligibility screening for clinical trials in a large community cancer center. *JCO clinical cancer informatics* **4**, 50–59 (2020).
18. Alexander, M. *et al.* Evaluation of an artificial intelligence clinical trial matching system in australian lung cancer patients. *JAMIA open* **3**, 209–215 (2020).
19. Nievas, M., Basu, A., Wang, Y. & Singh, H. Distilling large language models for matching patients to clinical trials. *J. Am. Med. Informatics Assoc.* **31**, 1953–1963 (2024).
20. Gupta, S. *et al.* Prism: Patient records interpretation for semantic clinical trial matching system using large language models. *NPJ digital medicine* **7**, 305 (2024).
21. Datta, S. *et al.* Autocriteria: a generalizable clinical trial eligibility criteria extraction system powered by large language models. *J. Am. Med. Informatics Assoc.* **31**, 375–385 (2024).

22. Beattie, J. *et al.* Utilizing large language models for enhanced clinical trial matching: a study on automation in patient screening. *Cureus* **16** (2024).
23. Bornet, A. *et al.* Analysis of eligibility criteria clusters based on large language models for clinical trial design. *J. Am. Med. Informatics Assoc.* **32**, 447–458 (2025).
24. Wornow, M. *et al.* Zero-shot clinical trial patient matching with llms. *NEJM AI* **2**, AIcs2400360 (2025).
25. Chen, H. *et al.* Enhancing patient-trial matching with large language models: a scoping review of emerging applications and approaches. *JCO Clin. Cancer Informatics* **9**, e2500071 (2025).
26. Kann, B. H., Hosny, A. & Aerts, H. J. Artificial intelligence for clinical oncology. *Cancer cell* **39**, 916–927 (2021).
27. Kehl, K. L. *et al.* Identifying oncology clinical trial candidates using artificial intelligence predictions of treatment change: a pilot implementation study. *JCO Precis. Oncol.* **8**, e2300507 (2024).
28. Nature Medicine. Show us the evidence for the value of medical AI. *Nat. Medicine* **32**, 1163, DOI: [10.1038/s41591-026-04389-4](https://doi.org/10.1038/s41591-026-04389-4) (2026). Editorial.
29. Gueguen, L. *et al.* A prospective pragmatic evaluation of automatic trial matching tools in a molecular tumor board. *npj Precis. Oncol.* **9**, 28 (2025).
30. Koopman, B. & Zuccon, G. A test collection for matching patients to clinical trials. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, 669–672 (2016).
31. Soboroff, I. Overview of trec 2021. In *TREC* (2021).
32. Roberts, K., Demner-Fushman, D., Voorhees, E. M., Bedrick, S. & Hersh, W. R. Overview of the trec 2022 clinical trials track. In *TREC* (2022).
33. Lowe, H. J. & Barnett, G. O. Understanding and using the medical subject headings (mesh) vocabulary to perform literature searches. *Jama* **271**, 1103–1108 (1994).
34. National Library of Medicine. ClinicalTrials.gov. <https://clinicaltrials.gov/> (2026).
35. National Institutes of Health Clinical Center. Search the studies. <https://clinicalstudies.info.nih.gov/> (2026).
36. OpenAI. Introducing gpt-5.5. <https://openai.com/index/introducing-gpt-5-5/> (2026).
37. Google. Gemini 3.5: frontier intelligence with action. <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-5/> (2026).
38. Liu, X. *et al.* Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the consort-ai extension. *The Lancet Digit. Heal.* **2**, e537–e548 (2020).
39. Rivera, S. C. *et al.* Guidelines for clinical trial protocols for interventions involving artificial intelligence: the spirit-ai extension. *The Lancet Digit. Heal.* **2**, e549–e560 (2020).
40. Syed, M. *et al.* Translating evidence into practice: adapting trialgpt for real-world clinical trial eligibility screening. *J. Am. Med. Informatics Assoc.* ocag006 (2026).
41. European Medicines Agency. Clinical trials information system. <https://www.ema.europa.eu/en/human-regulatory-overview/research-development/clinical-trials-human-medicines/clinical-trials-information-system>.

42. World Health Organization. International clinical trials registry platform. <https://www.who.int/tools/clinical-trials-registry-platform>.
43. Ministry of Health, Labour and Welfare, Japan. Japan registry of clinical trials. <https://jrct.mhlw.go.jp/en-top>.
44. Robertson, S. & Zaragoza, H. *The probabilistic relevance framework: BM25 and beyond*, vol. 4 (Now Publishers Inc, 2009).
45. Jin, Q. *et al.* Medcpt: Contrastive pre-trained transformers with large-scale pubmed search logs for zero-shot biomedical information retrieval. *Bioinformatics* **39**, btad651 (2023).
46. OpenAI. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/> (2025).
47. OpenAI. Introducing gpt-5.4. <https://openai.com/index/introducing-gpt-5-4/> (2026).
48. Google AI for Developers. Gemini 3.1 pro preview. <https://ai.google.dev/gemini-api/docs/models/gemini-3.1-pro-preview> (2026).
49. National Institutes of Health Office of Human Subjects Research Protections. Not human subjects research. <https://irbo.nih.gov/irb-review/not-human-subjects-research/>.
50. Schütze, H., Manning, C. D. & Raghavan, P. *Introduction to information retrieval*, vol. 39 (Cambridge University Press Cambridge, 2008).
51. Järvelin, K. & Kekäläinen, J. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Inf. Syst.* **20**, 422–446 (2002).
52. Voorhees, E. M. *et al.* The trec-8 question answering track report. In *TREC*, vol. 99, 77–82 (1999).

Acknowledgements

This research was supported by the Intramural Research Program of the National Institutes of Health (NIH). The contributions of the NIH author(s) are considered Works of the United States Government. This research was also partially supported by the NIH Pathway to Independence Award K99LM014903 (Q.J.). The findings and conclusions presented in this paper are those of the author(s) and do not necessarily reflect the views of the NIH or the U.S. Department of Health and Human Services.

We thank Juhi Gor and Aseem Aseem at the University of Illinois Cancer Center for data curation support for the retrospective and prospective Precision Oncology Tumor Board evaluations. We thank Zhizheng Wang at the National Library of Medicine, National Institutes of Health, for his valuable feedback on this work. We also thank the clinical reviewers, trial navigators, coordinators and recruitment staff at the participating sites for their contributions to clinical trial review workflows and feedback on TrialGPT 2.0 outputs.

NIH-TrialBench Consortium

National Cancer Institute (NCI): Stacey Doran, MD; Brooke Kania, DO; Charalampos S. Floudas, MD, DMSc, MS; Yoonhee Choi, MD; Aanika Warner, MD, MPH; Padma Rajagopal, MD, MPH, MSc; and Jung-Min Lee, MD.

National Human Genome Research Institute (NHGRI): Benjamin Solomon, MD; Christopher Ours, MD, MHS; and Precilla D’Souza, DNP, MSN, CRNP.

National Eye Institute (NEI): Mélanie Hébert, MD, MSc; Asmita Indurkar, MBBS, MS; Emily Chew, MD; and Tiarnán Keenan, MD, PhD.

National Heart, Lung, and Blood Institute (NHLBI): Emma Groarke, MD; and Gonca Ozcan, MD.

National Institute of Mental Health (NIMH): Mark Kvarta, MD, PhD.

National Institute of Dental and Craniofacial Research (NIDCR): Vivian Macdonald-Szymczuk, MD; Iris Hartley, MD; and Konstantinia Almpanti, DDS, MSc.

National Institute of Allergy and Infectious Diseases (NIAID): Elise O’Connell, MD; and Daniel Rogan, MD.

National Institute of Environmental Health Sciences (NIEHS): Stavros Garantziotis, MD.

National Institute on Drug Abuse (NIDA): Shiva Pandey, PA-C.

Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD): Christina Tatsi, MD, MHSc, PhD.

Author contributions

Y.F., Q.J. and Z.L. conceived the study. Y.F., Q.J. and Z.L. designed the TrialGPT 2.0 system evaluation. Y.F. developed the TrialGPT 2.0 system. S.T. implemented the TrialGPT 2.0 interface. L.H. coordinated the collection of synthetic vignettes from NIH Institutes and Centers. L.H. and M.G. provided input on clinical recruitment workflows. Y.F. performed the primary analyses, generated figures, and drafted the manuscript with input from Q.J., Z.W., J.S., and Z.L. N.N. and R.H.-T.N. led the UIC retrospective and prospective Precision Oncology Tumor Board evaluations. C.S.F. contributed to the design of the NCI retrospective evaluation and provided the 100 de-identified clinical vignettes. K.M. and C.M.M. performed the timed exhaustive-review annotations, and J.L.G. contributed clinical trial recruitment-workflow expertise. A.N., J.V., M. Bachini, L.R.-N. and K.R. contributed the CCF cohort and clinical review. N.C., M. Barnes and M.M. contributed the OPR cohort and recruitment-workflow expertise. D.G., C.E.G., H.S., and M. Burczynski contributed the UPMC cohort and clinical review. The NIH-TrialBench Consortium authored target-trial-informed synthetic vignettes and contributed clinical domain expertise for the NIH-TrialBench benchmark. Z.L. supervised the study. All authors interpreted the results, revised the manuscript, and approved the final version.

Competing interests

The authors declare no competing interests.

Supplementary Information

Supplementary Notes

- [Supplementary Note 1: Trial-fit factor coverage across matching policies](#)
- [Supplementary Note 2: Fit and confidence scores respond to clinical mismatch and missing information](#)
- [Supplementary Note 3: Clinician-comment analysis of non-recommended outputs](#)
- [Supplementary Note 4: Public benchmark comparison of TrialGPT 1.0 and TrialGPT 2.0](#)
- [Supplementary Note 5: Recommendation-generation time in the TrialGPT 2.0 interface](#)
- [Supplementary Note 6: Clinical reviewer acceptance of TrialGPT 2.0](#)

Supplementary Tables

- [Supplementary Table 1: Depth of trial-fit factor assessment by matching policy](#)
- [Supplementary Table 2: Clinical reviewer acceptance questionnaire](#)

Supplementary Figures

- [Supplementary Fig. 1: TrialGPT 2.0 score behavior under clinical mismatch and missing-information perturbations](#)
- [Supplementary Fig. 2: Clinician-comment analysis of discordant TrialGPT 2.0 Highly Recommended patient-trial pairs](#)
- [Supplementary Fig. 3: Public benchmark comparison of TrialGPT 2.0 and TrialGPT 1.0 across LLM backbones](#)
- [Supplementary Fig. 4: TrialGPT 2.0 web-based interface for trial recommendation review](#)
- [Supplementary Fig. 5: TrialGPT 2.0 recommendation-generation runtime across evaluation sources](#)
- [Supplementary Fig. 6: Clinical reviewer acceptance responses for TrialGPT 2.0](#)
- [Supplementary Fig. 7: NIH-TrialBench target-trial recovery](#)

Supplementary Note 1: Trial-fit factor coverage across matching policies

The trial-fit factors shown in Fig. 1 represent factors that TrialGPT 2.0 can assess rather than a fixed checklist applied identically in every workflow. Each site- or evaluation-specific matching policy determines which factors are represented and the depth with which they are assessed. Factor coverage was categorized as follows:

- **Dedicated:** the matching policy contains specific instructions for assessing the factor.
- **Partial:** the factor is considered only in certain circumstances, through the raw trial eligibility criteria or through instructions related to another factor, rather than through separate factor-specific instructions.
- **Triage:** the factor is used for high-level screening based on brief referral information rather than comprehensive assessment.

These categories describe how each factor is represented in the matching policy. All matching policies receive the patient context and raw trial eligibility criteria; therefore, a factor categorized as Partial may still influence the assessment when relevant information is available. Therapeutic benefit was not categorized separately because it was treated as a composite clinical judgment reflected in the overall fit assessment and rationale rather than as a standalone factor-specific component.

Diagnosis or histology. All matching policies assess diagnosis or disease subtype. Histology is considered when specified in the raw trial eligibility criteria. UIC, CCF, UPMC, NCI, and NIH-TrialBench contain dedicated diagnosis-matching instructions, whereas OPR uses diagnosis or phenotype primarily for referral triage.

Disease status. UIC and CCF assess disease and stage fit through dedicated instructions, including factors such as advanced or metastatic disease, measurable disease and treatment-naïve versus previously treated settings. UPMC explicitly considers selected disease-status characteristics, including measurable disease, recurrence or second-primary status and prior curative treatment, but does not contain a comprehensive disease-status component. NCI, OPR and NIH-TrialBench consider disease fit through the patient context, trial criteria and broader relevance assessment, without a dedicated component covering metastatic, recurrent, stage or no-evidence-of-disease status.

Biomarker or pathway match. UIC, CCF, NCI and NIH-TrialBench contain explicit instructions for assessing required biomarkers, biomarker-defined disease and biomarker-specific cohorts. Pathway relevance is also emphasized in UIC and CCF but is not separately represented in the NCI and NIH-TrialBench policy. UPMC considers biomarkers and genomic categories primarily in relation to cohort or arm assignment. OPR treats the absence of a required biomarker as an important screening concern but does not emphasize broader pathway matching.

Line of therapy. UIC and CCF contain dedicated line-of-therapy components that compare the patient's prior treatment-line count with trial-level or cohort-level requirements. UPMC, NCI, NIH-TrialBench and OPR consider line of therapy when relevant information is present in the patient context or raw trial criteria, but do not contain dedicated line-counting components.

Prior treatment. UIC, CCF and UPMC explicitly assess prior systemic therapy, treatment timing, washout requirements and prohibited or concurrent therapies. NCI, OPR and NIH-TrialBench focus primarily on prohibited prior therapies or prior-treatment information directly relevant to the trial criteria.

Supplementary Table 1. Depth of trial-fit factor assessment by matching policy. Symbols indicate whether each factor is addressed through specific policy instructions, considered partially, or applied at a triage level.

Trial-fit factor	OPR	CCF	UIC	UPMC	NCI	NIH-TrialBench
Diagnosis or histology	△	●	●	●	●	●
Disease status	○	●	●	○	○	○
Biomarker or pathway match	○	●	●	○	●	●
Line of therapy	○	●	●	○	○	○
Prior treatment	○	●	●	●	○	○
Cohort or arm applicability	○	●	●	●	●	●
Inclusion and exclusion criteria	△	●	●	●	●	●
Key: ● = Dedicated ○ = Partial △ = Triage						

Cohort or arm applicability. UIC, CCF, UPMC, NCI and NIH-TrialBench contain explicit logic for evaluating umbrella or master protocols and determining whether the patient fits at least one relevant cohort or arm. OPR considers alignment with the referral goal and broad cohort preferences but does not perform comprehensive multi-arm assignment.

Inclusion and exclusion criteria. UIC, CCF, UPMC, NCI and NIH-TrialBench use the raw trial inclusion and exclusion criteria as the primary basis for formal eligibility assessment. OPR applies these criteria in a triage-oriented manner, emphasizing clear hard exclusions, demographic mismatches, numeric-threshold violations and alignment with the referral goal.


Supplementary Note 2: Fit and confidence scores respond to clinical mismatch and missing information

We performed complementary analyses to evaluate whether the TrialGPT 2.0 fit and confidence scores behave as intended. We first tested whether the fit score distinguishes patient-specific compatibility from partial or nonspecific relevance. We then tested whether it responds to isolated incompatibilities in individual trial-fit factors. Finally, we examined whether the confidence score responds selectively to the removal of information needed to assess a candidate trial.

Fit-score discrimination using partial-match and patient-shuffled controls

For the fit-score comparison, we used existing TrialGPT 2.0 matching outputs from NIH-TrialBench and the NCI, UIC and CCF retrospective cohorts. These evaluation sources were selected because their matching policies contained specific instructions for both diagnosis or disease matching and biomarker matching, the two dimensions used to construct the partial-match controls (Supplementary Note 1 and Supplementary Table 1). Clinician-recommended pairs were defined as patient-trial pairs that TrialGPT 2.0 categorized as Highly Recommended and clinicians labeled as Recommended. We compared these pairs with three control groups designed to preserve apparent relevance while disrupting the patient-specific match (Supplementary Fig. 1a).

The first two groups were partial-match controls selected from each patient's scored TrialGPT 2.0 output. Disease-matched/biomarker-mismatched controls were patient-trial pairs in which the trial appeared relevant to the patient's disease but required a different or undocumented biomarker or pathway. Biomarker-matched/disease-mismatched controls were pairs in which the patient and trial shared a relevant biomarker or pathway, but the trial disease context did not



Supplementary Fig. 1. TrialGPT 2.0 score behavior under clinical mismatch and missing-information perturbations. **a**, Fit-score discrimination among clinician-recommended patient-trial pairs, disease-matched pairs with biomarker mismatch, biomarker-matched pairs with disease mismatch and shuffled clinician-recommended patient-trial pairs. P values were obtained using two-sided Welch's t -tests comparing patient-level mean fit scores for clinician-recommended trials versus each control condition. **b**, Confidence-score change after removing trial-irrelevant or trial-critical information from patient summaries. P values were obtained using two-sided paired t -tests comparing the original confidence score with the confidence score after information removal for the same patient-trial pair. **c**, Fit-score sensitivity to counterfactual one-factor hard mismatches in diagnosis or histology, stage or disease status, biomarker or pathway, line of therapy, prior treatment, cohort or arm applicability, inclusion criteria and exclusion criteria. Gray lines connect paired original and perturbed scores for the same patient-trial anchor. P values were obtained using two-sided paired t -tests comparing the original fit score with the perturbed fit score for the same patient-trial anchor. Dashed horizontal lines indicate the Highly Recommended threshold in **a** and **c** and no confidence-score change in **b**. Boxes show the interquartile range, center lines show medians and whiskers extend to the most extreme values within $1.5\times$ the interquartile range. Points indicate individual patient-trial pairs in **a**, paired perturbation examples in **b** and paired original or perturbed observations in **c**.

match the patient.

Candidate partial-match pairs were reviewed using GPT-5.4⁴⁷, with the patient summary, trial criteria and TrialGPT 2.0 explanation provided as inputs. The TrialGPT 2.0 explanation was included to clarify the assessed rationale and support review of ambiguous multi-arm or basket-

trial contexts. A candidate was retained only when the assigned mismatch category was confirmed and no plausible applicable cohort or arm was identified.

The third group consisted of patient-shuffled controls. Starting from a clinician-recommended pair comprising patient A and trial X, we identified other patients for whom trial X also appeared in the scored TrialGPT 2.0 output. Each resulting pair between trial X and a different patient was retained as a shuffled control, using the fit score already assigned by TrialGPT 2.0 to that patient-trial pair.

Fit scores were aggregated to patient-level means, and clinician-recommended pairs were compared with each control condition using two-sided Welch's *t*-tests. Clinician-recommended pairs received consistently high fit scores, whereas both partial-match controls scored substantially lower. These results indicate that the fit score does not rely on disease or biomarker similarity alone. Patient-shuffled controls also scored lower, consistent with the fit score reflecting the specific patient-trial pairing rather than trial-level attractiveness (Supplementary Fig. 1a).

Fit-score sensitivity to counterfactual hard mismatches

We next tested whether the fit score responds to individual trial-fit factors (Fig. 1b). We used clinician-recommended patient-trial pairs from NIH-TrialBench and the retrospective cohorts as positive anchors. For each anchor, GPT-5.4 generated a minimally edited patient profile that disrupted one prespecified trial-fit factor while keeping the trial text fixed and preserving the remainder of the patient profile whenever possible. For example, a biomarker perturbation changed a required positive biomarker to negative or mutually exclusive, whereas a stage perturbation changed the disease status to one incompatible with the trial requirement (Supplementary Fig. 1c).

Generated perturbations were retained only after quality-control review. A retained perturbation had to satisfy five criteria: the original patient-trial pair was clinician-recommended; the targeted factor was relevant to the candidate trial, as indicated by the raw trial eligibility criteria or the corresponding matching policy; the edited patient profile was internally consistent; the edit targeted the specified factor without introducing unrelated changes; and the edited patient was clearly incompatible with the trial for that factor. Multi-arm and basket trials were reviewed carefully, and perturbations were excluded when another arm or cohort remained plausibly applicable.

We did not perturb therapeutic relevance or benefit separately because it represents a composite clinical judgment rather than a single objective trial-fit variable that can be minimally changed without altering multiple clinical facts.

TrialGPT 2.0 was rerun on each original and counterfactual patient-trial pair using the same trial and site-specific prompt context. For Supplementary Fig. 1c, we used a balanced set of 50 quality-controlled perturbations for each factor group: diagnosis or histology, stage or disease status, biomarker or pathway, line of therapy, prior treatment, cohort or arm applicability, inclusion criteria and exclusion criteria. Each counterfactual fit score was compared with the original fit score for the corresponding clinician-recommended anchor pair using a two-sided paired *t*-test.

Single-factor hard mismatches produced marked reductions in fit score across all evaluated domains. These results show that the fit score responds to clinically important incompatibilities even when the trial and the remainder of the patient profile are held constant (Supplementary Fig. 1c).

Confidence-score sensitivity to missing information

The confidence score addresses a different property: whether the available patient information is sufficient to support the fit judgment. To evaluate this behavior, we used 100 original patient-trial anchor pairs from NIH-TrialBench and the retrospective cohorts. For each pair, GPT-5.4 generated two minimally edited patient profiles while keeping the candidate trial fixed: one with trial-irrelevant information removed and one with trial-critical information removed. This yielded 100 trial-irrelevant information-removal pairs and 100 trial-critical information-removal pairs (Supplementary Fig. 1b).

Trial-irrelevant information comprised patient details that were not expected to affect assessment of the candidate trial, such as unrelated medical history, non-decisive comorbidities or details not referenced by the trial criteria. Trial-critical information comprised evidence needed to assess eligibility or clinical fit, such as biomarker status, disease stage or status, line of therapy, prior treatment exposure, ECOG performance status, organ function or central nervous system involvement.

Generated perturbations were retained only after quality-control review. The removed information had to be present in the original patient profile, the edited profile had to remain internally consistent, and the edit had to remove information rather than replace it with an incompatible value. This distinction separated missing-information perturbations from the counterfactual hard mismatches described above: missing or obscured evidence was expected to affect confidence, whereas direct incompatibility was expected to affect fit.


TrialGPT 2.0 was rerun on the same candidate trial using the perturbed patient context and the same site-specific prompt context. For each pair, we calculated the confidence-score change relative to the original assessment. Changes were summarized separately for trial-irrelevant and trial-critical information removal, and the confidence score under each removal condition was compared with the corresponding original score using a two-sided paired *t*-test.

Removing trial-irrelevant information produced only small confidence changes, whereas removing trial-critical information led to larger reductions. These results support the intended role of the confidence score in reflecting whether key information needed for patient-trial assessment is available (Supplementary Fig. 1b).

Supplementary Note 3: Clinician-comment analysis of non-recommended outputs

To understand why some top-ranked TrialGPT 2.0 recommendations were not judged clinically appropriate, we analyzed free-text clinician comments attached to patient-trial pairs that TrialGPT 2.0 categorized as Highly Recommended but clinicians labeled as Eligible but Not Recommended or Ineligible.

Each comment was assigned one primary reason for concern according to the main issue expressed by the clinician. Categories comprised condition or biomarker fit issue, limited clinical benefit, standard-of-care or treatment-sequence conflict, trial-intent issue, protocol issue and other. Condition or biomarker fit issues included mismatches in disease, histology, molecular subtype or required biomarker. Limited clinical benefit included cases in which a trial was technically possible but was judged to offer low value because of expected benefit, toxicity, patient goals or better available options. Standard-of-care or treatment-sequence conflicts included concerns related to treatment timing, prior-therapy requirements or line-of-therapy logic. Trial-intent issues included studies whose purpose did not match the patient's clinical need, such as imaging,




Supplementary Fig. 2. Clinician-comment analysis of discordant TrialGPT 2.0 Highly Recommended patient-trial pairs. The alluvial plot summarizes patient-trial pairs assigned to the Highly Recommended category by TrialGPT 2.0 but labeled by clinicians as Eligible but Not Recommended or Ineligible. The left side shows clinician annotations, and the right side shows the main reason for concern identified from clinician comments, including condition or biomarker fit issues, limited clinical benefit, standard-of-care (SOC) or treatment-sequence conflicts, trial-intent issues, protocol issues and other concerns. Counts and percentages are shown for each annotation and concern category, with representative clinician comments displayed as examples.

screening, registry, prevention or supportive-care studies when therapeutic treatment was sought. Protocol issues included specific eligibility criteria, required protocol participation, cohort rules, enrollment requirements or trial-status concerns. Comments that did not fit these categories or were too nonspecific to assign confidently were labeled as other.

Category assignment was performed at the individual-comment level. When a comment contained multiple concerns, we assigned the category corresponding to the dominant reason that the clinician rejected or deprioritized the trial. Assignments were manually reviewed for consistency, with particular attention to overlapping themes such as limited benefit versus trial intent and biomarker mismatch versus cohort or protocol mismatch. Comments that did not contain a substantive clinical concern were not assigned to a reason-for-concern category.

Reason-for-concern categories were summarized overall and stratified by clinician annotation, comparing Eligible but Not Recommended with Ineligible cases. Clinician comments identified both patient-trial compatibility concerns, including condition, biomarker and protocol mismatches, and clinical-priority concerns, including limited expected benefit, treatment-sequence conflicts and mismatch between trial intent and the patient's care goal. The presence of many Eligible but Not Recommended annotations indicates that discordant outputs were not limited to binary eligibility failures, but frequently reflected clinical judgment about trial appropriateness and priority (Supplementary Fig. 2).

Supplementary Note 4: Public benchmark comparison of TrialGPT 1.0 and TrialGPT 2.0



Supplementary Fig. 3. Public benchmark comparison of TrialGPT 2.0 and TrialGPT 1.0 across LLM backbones. **a**, TrialGPT 2.0 replaces the TrialGPT 1.0 criterion-level aggregation workflow with criterion-aware trial-level scoring. **b**, Computational efficiency across SIGIR, TREC2021 and TREC2022 using GPT-4.1⁴⁶, GPT-5.4⁴⁷ and Gemini-3.1-pro-preview⁴⁸, measured by matching speed and token use per patient-trial pair. Percentages indicate relative changes for TrialGPT 2.0 versus TrialGPT 1.0. **c**, Eligibility-ranking performance across public benchmarks and LLM backbones, evaluated using Precision@10, normalized discounted cumulative gain at 10 (nDCG@10), mean reciprocal rank (MRR) and mean average precision (MAP). Bars denote mean performance. For individual benchmarks, error bars show approximate 95% confidence intervals (mean $\pm$ 1.96 $\times$ standard error across patient queries); for Avg bars, error bars summarize variation across the three benchmark-level means. Percentage annotations indicate the relative change of TrialGPT 2.0 compared with TrialGPT 1.0.

We evaluate TrialGPT 2.0 on established public patient-trial matching benchmarks, including the 2016 Special Interest Group on Information Retrieval (SIGIR) test collection³⁰ and the 2021 and 2022 Clinical Trials tracks of the Text REtrieval Conference (TREC)^{31,32}.

We compare TrialGPT 2.0 with the original TrialGPT framework¹⁰, referred to here as TrialGPT 1.0. TrialGPT 1.0 performs criterion-level matching: it decomposes raw eligibility criteria into inclusion and exclusion criteria, evaluates each criterion against the patient record and aggregates criterion-level outputs into a trial-level score and explanation. TrialGPT 2.0 retains criterion-aware assessment but generates trial-level judgments directly, reviewing trial criteria jointly and returning a fit score, confidence estimate, rank and structured explanation for each trial (Supplementary Fig. 3a).

We assess ranking performance using Precision@10⁵⁰, nDCG@10⁵¹, mean reciprocal rank (MRR)⁵² and mean average precision (MAP)⁵⁰, standard information-retrieval metrics that capture complementary aspects of ranked retrieval quality, including top-ranked precision, rank-sensitive relevance, early retrieval of relevant trials and ranking quality across relevant trials. Across SIGIR, TREC 2021 and TREC 2022, TrialGPT 2.0 improves all evaluated ranking metrics relative to TrialGPT 1.0. Averaged across benchmarks, TrialGPT 2.0 increases Precision@10 by 4.5%, nDCG@10 by 5.7%, MRR by 9.8% and MAP by 10.2% (Supplementary Fig. 3c). These gains indicate that joint trial-level assessment improves benchmark ranking quality, a prerequisite for using TrialGPT 2.0 to prioritize candidate trials for clinician review.

TrialGPT 2.0 also reduces inference burden. Averaged across benchmarks, processing speed increases by 3.20-fold, while input and output tokens per patient-trial pair decrease by 58% and 73%, respectively (Supplementary Fig. 3b). This efficiency gain is consistent with the architectural change from criterion-level aggregation to joint trial-level assessment: criterion-level matching requires repeated assessment of individual eligibility criteria and longer intermediate outputs, whereas TrialGPT 2.0 produces a compact structured assessment for each patient-trial pair. In the web-based implementation, candidate-trial assessments run in parallel using multiprocessing, allowing the interface to display ranked matching results for one case in approximately 22 seconds on average.

Additional analyses with GPT-5.4⁴⁷ and Gemini-3.1-pro-preview⁴⁸ show a similar overall pattern of improved efficiency and favorable average ranking performance with TrialGPT 2.0 (Supplementary Fig. 3). We use GPT-4.1 in the web-based implementation because it provides a favorable trade-off between ranking performance and computational efficiency.

Supplementary Note 5: Recommendation-generation time in the TrialGPT 2.0 interface

To assess the operational feasibility of TrialGPT 2.0 in the interface setting (Supplementary Fig. 4), we measured recommendation-generation time across five implementation settings: CCF, NCI, OPR, UIC and UPMC. For each source, we sampled five patient cases and ran each case three times using the corresponding site-specific trial corpus and matching configuration. Runs were performed in the interface runtime environment using 128 worker processes.

The timing experiment used the same trial-selection configuration as the interface. Retrieval was used only for UIC, where the trial corpus was first filtered using the hybrid retrieval strategy based on BM25 and MedCPT; by default, the top 267 retrieved trials were passed to the TrialGPT 2.0 backend for detailed assessment. For CCF, NCI, OPR and UPMC, retrieval was not used and TrialGPT 2.0 scored the full site-specific trial corpus directly. For UIC, retrieval was applied before backend assessment.

Timing was measured from initiation of the matching process to completion of the TrialGPT 2.0 recommendation-generation pipeline. The measurement included trial retrieval when applicable,

National Library of Medicine

[An official website of the United States government](#)
[Here's how you know](#)

[HOME](#)
[NEWS](#)
[DEMO](#)
[FAQ](#)
[CODE & DATA](#)
[ABOUT](#)

[Log in](#)

Filter & Scoring

Recommendation

- ☐ Highly Recommended
- ☐ Possible Match
- ☐ Low Fit

Phase

- ☐ Phase I
- ☐ Phase II
- ☐ Phase III
- ☐ Phase IV

Scope

- ☐ General
- ☐ Biomarker-specific

Sort By

- ☒ Score
- ☐ Distance

Location

Enter one or more cities or ZIP codes, separated by a semicolon (;).

[Apply Filters](#)

[Reset Filters](#)

Matched Trials [52]

[New Matching](#)

[PDF](#)

Highly Recommended (21)

Possible Match (31)

Low Fit (0)

Interventional Trials

Non-Interventional Trials

1
Molecular Profiling and Targeted Therapy for Advanced Non-Small Cell Lung Cancer, Small Cell Lung Cancer, and Thymic Malignancies
NCT01306045
Phase II • Biomarker-specific • Completed
Eligible Reasons: 1. Diagnosis of metastatic non-small cell lung cancer (NSCLC), which is explicitly included in trial eligibility...
View details

2
Metarrestin (ML-246) in Subjects With Metastatic Solid Tumors
NCT04222413
Phase I • General • Recruiting
Eligible Reasons: 1. The patient is an adult (age 61) with a histologically confirmed metastatic solid tumor (NSCLC), which fits the Phase IA inclusion for adults with solid tumors...
View details


3
Targeted Therapy Directed by Genetic Testing in Treating Patients With Locally Advanced or Advanced Solid Tumors, The ComboMATCH Screening Trial
NCT05564377
Phase II • General • Recruiting
Eligible Reasons: 1. Diagnosis of metastatic NSCLC (a solid tumor) matches the trial's inclusion of locally advanced or advanced solid tumors...
View details

4
A Study of a Selective T Cell Receptor (TCR) Targeting, Bifunctional Antibody-fusion Molecule STAR0602 in Participants With Advanced Solid Tumors
NCT05592626
Phase I/II • Biomarker-specific • Recruiting
Eligible Reasons: 1. Diagnosis of metastatic non-small cell lung cancer (NSCLC), which is explicitly listed as an eligible tumor type for Phase 2...
View details

Supplementary Fig. 4. TrialGPT 2.0 web-based interface for trial recommendation review. The interface accepts a patient clinical summary and returns matched trials grouped by TrialGPT 2.0 recommendation category. Reviewers can filter and sort recommendations, export matched results and inspect trial-level explanations. Each trial card presents trial metadata and a structured rationale, including eligibility-supporting evidence, missing information and ineligibility concerns.

backend assessment, output parsing, result writing and score-sorted PDF generation. Manual clinician review, downstream chart review and coordinator follow-up were not included.

For each case, we averaged the three repeated runs to obtain a case-level runtime estimate. We then summarized recommendation-generation time across cases within each source using the mean and standard deviation. Across all 25 sampled cases, TrialGPT 2.0 generated recommendations in 21.7 s on average, supporting use of the interface in trial-review workflows that require short turnaround time for candidate-trial generation (Supplementary Fig. 5).



Supplementary Fig. 5. TrialGPT 2.0 recommendation-generation runtime across evaluation sources. Five cases are sampled from each evaluation source, and each case is run three times using 128 worker processes. The three runs are averaged to obtain a case-level runtime estimate. Bars show the mean runtime across cases within each source, error bars show the standard deviation across cases and points indicate case-level means; the Avg bar pools case-level means across all sources. Runtime includes trial retrieval when applicable, backend assessment, output parsing, result writing and generation of score-sorted PDF reports.

Supplementary Note 6: Clinical reviewer acceptance of TrialGPT 2.0

Supplementary Table 2. Clinical reviewer acceptance questionnaire.


Instructions: Please indicate your level of agreement with each statement based on your experience using TrialGPT 2.0 for clinical trial review.

5-point Likert scale: 1 = Strongly disagree, 2 = Disagree, 3 = Neutral, 4 = Agree, 5 = Strongly agree.

No.	Statement	Level of agreement (select one)				
		1 Strongly disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly agree
1	Usability: TrialGPT 2.0 was user-friendly.	☐	☐	☐	☐	☐
2	Workflow fit: TrialGPT 2.0 fit well into my clinical workflow.	☐	☐	☐	☐	☐
3	Efficiency: TrialGPT 2.0 reduced the time required for trial review.	☐	☐	☐	☐	☐
4	Clarity: TrialGPT 2.0 outputs were easy to understand.	☐	☐	☐	☐	☐
5	Decision confidence: TrialGPT 2.0 increased my confidence in my trial-screening decisions.	☐	☐	☐	☐	☐
6	Usefulness: TrialGPT 2.0 outputs were relevant and meaningful for trial screening.	☐	☐	☐	☐	☐
7	Content safety: I did not observe violent, offensive, or otherwise harmful content in the TrialGPT 2.0 outputs I reviewed.	☐	☐	☐	☐	☐
8	Satisfaction: Overall, I was satisfied with TrialGPT 2.0.	☐	☐	☐	☐	☐


Clinical reviewer acceptance of TrialGPT 2.0 was assessed using an 8-item questionnaire after reviewers used or reviewed TrialGPT 2.0 outputs in clinical trial matching workflows. The questionnaire covered usability, workflow fit, perceived efficiency, clarity, decision confidence, usefulness, content safety and overall satisfaction. Each item used a 5-point Likert scale, where 1 indicated strongly disagree and 5 indicated strongly agree (Supplementary Table 2).

Thirteen clinical reviewers completed the questionnaire. Across 104 item-level ratings, 85 ratings were favorable, defined as a score of 4 or 5. Mean item scores ranged from 3.85 to 4.92. Content safety received the highest rating, with a mean score of 4.92 and favorable responses from all



Supplementary Fig. 6. Clinical reviewer acceptance responses for TrialGPT 2.0. Clinical roles of questionnaire respondents and response distributions across the eight reviewer-acceptance items. The questionnaire was completed by 13 clinical reviewers after use or review of TrialGPT 2.0 outputs in clinical trial matching workflows. Items assessed usability, workflow fit, perceived efficiency, clarity, decision confidence, usefulness, content safety, and overall satisfaction on a 5-point Likert scale, where 1 indicates strongly disagree, and 5 indicates strongly agree. Bar plots show the number of respondents selecting each rating for each item; mean annotations indicate the average rating for each item.

reviewers. Overall satisfaction and usability each received a mean score of 4.54, followed by clarity and workflow fit, with mean scores of 4.38 and 4.31, respectively. Ratings were more variable for perceived efficiency and decision confidence, with mean scores of 3.85 and 3.92, respectively. These results suggest that reviewers generally found TrialGPT 2.0 safe, usable, understandable, and satisfactory, while perceived time savings and confidence gains may depend on workflow context, reviewer role, and continued user calibration (Supplementary Fig. 6).



Supplementary Fig. 7. NIH-TrialBench target-trial recovery. Each synthetic vignette is associated with a prespecified target trial. For each vignette, we record the TrialGPT 2.0 recommendation category assigned to the target trial. The table and stacked bars show the number and proportion of target trials categorized as Highly Recommended, Possible Match, or Low Fit, separately for each contributing NIH Institute or Center and overall.