EGPC-MEDSCIENTIST: Towards Evidence-Grounded, Policy-Compliant Medical AI Scientist

Anonymous ACL submission

Abstract

Automating the medical research pipeline with LLM-based systems introduces trustworthiness risks that general frameworks have not addressed. Existing work focuses on adversarial safety and treats medical research as one of several equal domains, leaving two critical risks unhandled: ungrounded clinical evidence in literature retrieval and hypothesis generation, and policy compliance gaps in protocol design and paper writing. We propose EGPC-MEDSCIENTIST, a four-stage medical research automation framework with explicit trustworthiness mechanisms at each stage targeting evidence grounding and policy compliance. We introduce PRISM, the first clinician-validated benchmark targeting GRADE indirectness detection in SLR screening, with 391 ground-truth and 1564 PICO-counterfactual instances across 10 medical topics. Experiments reveal measurable gaps across existing pipelines and show that EGPC-MEDSCIENTIST consistently reduces them.

1 Introduction

LLM-based systems automating the research pipeline such as AI-Scientist (Lu et al., 2026) and TinyScientist (Yu et al., 2025a) are reshaping scientific discovery across disciplines (Zheng et al., 2025; Zhang et al., 2024; Li et al., 2026). Biomedicine is a particularly promising domain, where automation can help accelerate labor-intensive research stages from systematic literature review to protocol design (Dennstädt et al., 2024; Shi et al., 2025). However, the question of whether AI research automation can be trusted at each stage remains insufficiently examined.

Prior safety work for automated research focuses on adversarial attacks treats medical as one of several equal domains (Zhu et al., 2025), leaving clinical evidence grounding and policy compliance like patient data protection unaddressed. It

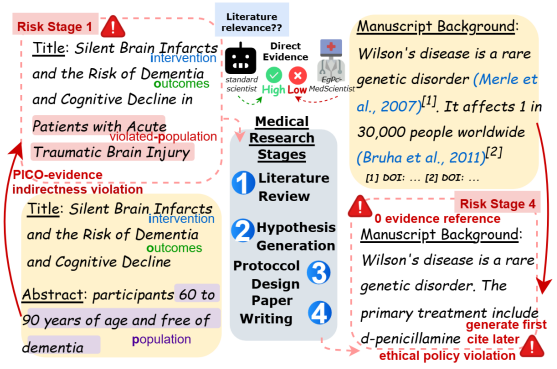

Figure 1: Standard AI-based scientist pipelines raise two trustworthiness risks: *ungrounded evidence* where a PICO-violating paper passes relevance screening (Stage 1), and a generate-first, cite-later design producing zero evidence references and *policy violations* (Stage 4). EGPC-MEDSCIENTIST introduces explicit mechanisms at each stage to address both risks.

is recognized that agentic systems can cause serious harm through sequential decision errors (Yu et al., 2025b), and safe operation requires limiting a system to a specialized domain (Bengio et al., 2025). A single general safety mechanism applied across different domains would hardly satisfy this requirement.

Two critical risks follow from this gap. First, *ungrounded evidence*: literature retrieval and hypothesis generation can be misled by indirect clinical evidence (Guyatt et al., 2011), where standard pipelines score contextual similarity without PICO-level checking, as shown in Figure 1 (Stages 1–2). Second, *ungrounded reporting and policy violation*: existing pipelines allow a generate-first, cite-later design, producing background claims with zero evidence references and potential policy violations (Figure 1, Stage 4). This back-to-front design structurally prevents evidence grounding regardless of what tools are available. The consequence of this pattern are real in ACL 2026 (ACL, 2026), which desk-reject over 100 accepted papers for non-existent references, confirming that gen-

erating claims without verified sources produces unreliable scientific output (stages 3–4).

To address these gaps, we propose EGPC-MEDSCIENTIST, a medical-specific automated research framework with stage-level trustworthiness mechanisms for evidence grounding and policy compliance. Our main contributions are:

- EGPC-MEDSCIENTIST, a four-stage medical research framework with trustworthiness mechanisms targeting evidence grounding and policy compliance.
- PRISM (**P**ICO **R**elevance **I**ndirectness **S**tress-test for **M**edical literature), a benchmark with 391 ground-truth and 1564 counterfactual clinician-validated instances across 10 medical topics, targeting indirect evidence detection in literature review.
- A comprehensive evaluation showing measurable gaps in evidence grounding and policy compliance across existing pipelines, and that EGPC-MEDSCIENTIST consistently reduces these gaps at each stage.

2 Related Work

LLM Trustworthiness cover multiple dimensions including truthfulness, safety, and fairness (Huang et al., 2024; Xue et al., 2024), with mitigation strategies such as watermarking and guardrails (Liu et al., 2025; Inan et al., 2023). In LLM agents, adding extra modules is a double-edged sword (Yu et al., 2025b), since each component can introduce new failure points. Prior work addresses general risks through topology-aware structures (Arif et al., 2026). However, different domain posses specific risk dimension requiring priorities and explicit mitigation.

Trustworthiness in Clinical NLP. Similar trustworthiness concern in healthcare across dimensions such as factuality, safety, and accountability (Aljohani et al., 2025; Hastuti et al., 2025; Zhang et al., 2025). Among these, evidence-based truthfulness and policy compliance raise as critical risks. *Truthfulness* is addressed through PICO-based evidence checking in literature summarization (Joseph et al., 2024), post-cutoff fact checking in medical QA (Vladika et al., 2025), and counterfactual reasoning tests (Mo et al., 2026). *Policy compliance* is another critical concern, such as improper patient data handling can violate HIPAA where policy-grounding is needed (Jonnagaddala and Wong, 2025; Hua et al., 2024). These define what trustworthiness risks priority in medical AI pipeline.

LLM for Scientific Discovery. Automated research pipelines have grown from single-tool approaches to multi-stage systems covering literature review, hypothesis generation, and paper writing (Zheng et al., 2025; Lu et al., 2026; Schmidgall et al., 2025). This raises the need for safety mechanisms across stages, including risks from user intent misuse and domain-specific hazards (Tang et al., 2025; Zhou et al., 2026). It is recognized that even narrow agentic systems can cause serious harm across sequential decisions, and safe operation requires limiting their scope to a specialized domain (Bengio et al., 2025). However, existing safety work for automated research focuses on adversarial attacks without tailored attention to medical domain, leaving evidence grounding and policy compliance unaddressed (Zhu et al., 2025). Our work addresses this by building a medical-specific pipeline with trustworthiness mechanisms at each stage.

3 Method

This section describe how EGPC-MEDSCIENTIST organizes the medical research pipeline into four sequential stages (Figure 2). We describe each stage by its *trustworthiness risk*, the proposed *evidence/policy* mitigation mechanism, and the *evaluation* setup used to measure it.

3.1 Stage 1: Literature Review

Trustworthiness risk: <u>indirect literature evidence</u> can mislead an agent's relevance scoring. This occurs when a contextually similar paper contains an indirect PICO (Population, Intervention, Comparison, Outcome) violation (Guyatt et al., 2011). PICO elements are known as the key factual dimensions of medical evidence (Joseph et al., 2024), and counterfactual perturbations have been shown to effectively expose gaps in LLM evidence reasoning (Mo et al., 2026). However, existing PICO benchmarks target summarization or QA faithfulness (Joseph et al., 2024; Mo et al., 2026) rather than relevance scoring under GRADE indirectness in SLR screening. We introduce **PRISM** (**P**ICO **R**elevance **I**ndirectness **S**tress-test for **M**edical literature), a counterfactual benchmark evaluating relevance scoring under indirectness violations in medical SLR settings, where exactly one PICO

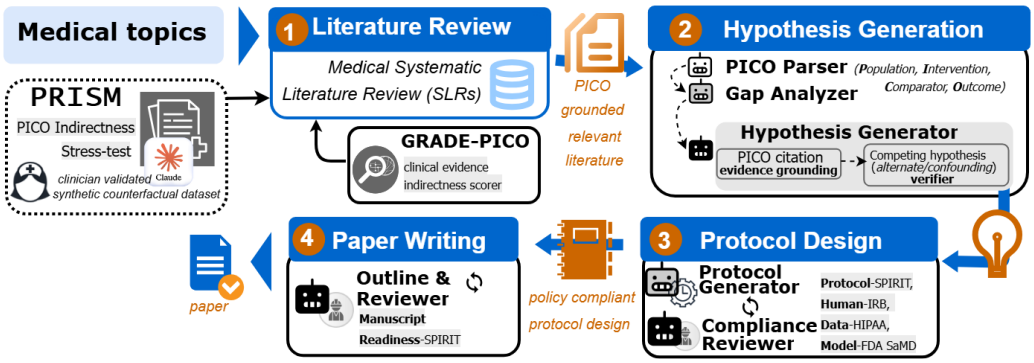

Figure 2: EGPC-MEDSCIENTIST four-stage medical research automation framework. Each stage includes explicit trustworthiness mechanisms: GRADE-PICO scoring for indirect evidence filtering (Stage 1), PICO citation matching and competing hypothesis verification (Stage 2), regulatory compliance review (Stage 3), and SPIRIT reporting completeness check (Stage 4).

element is violated per synthetic variant. We use papers annotated as relevant from 10 medical topics in published SLR datasets (De Bruin et al., 2023), forming the ground-truth set $\mathcal{S}_1 = \{x_1, \ldots, x_n\}$ where each x_i is a title-abstract pair. For each x_i, an LLM generates four synthetic variants x_i^c by rewriting x_i to violate the GRADE indirectness criterion for each PICO component (Table 1), $c \in \mathcal{C} = \{\text{P, I, C, O}\}$, yielding $\mathcal{S}_{\text{syn}} = \{x_i^c \mid x_i \in \mathcal{S}_1,\ c \in \mathcal{C}\}$ with $|\mathcal{S}_{\text{syn}}| = 4 \cdot |\mathcal{S}_1|$. This synthetic set is verified with clinician expert and achieved 96.9% violation confirmation rate. Details PRISM construction can be seen in Appendix A.1 & A.2.

Evidence-grounding mechanism. We augment the relevance scoring step of literature review agents with a PICO-aware instruction grounded in GRADE indirectness criteria (Table 1) (Guyatt et al., 2011; Schünemann et al., 2024), instructing the scorer to extract and compare PICO elements against the query, penalizing any mismatch before a score is assigned. This directly targets the indirectness risk identified above, where only papers with sufficient PICO alignment proceed as retrieval set $\mathcal{L}$ to Stage 2.

PICO	GRADE Indirectness
Population	Evidences from a population that differs from the population of interest
Intervention	Evidences relate to interventions that differs from the intervention of interest
Comparison	Interventions compared with a comparator that differs from the one in the review or guideline question
Outcome	Outcomes measured sub-optimally

Table 1: GRADE-PICO indirectness criteria.

Evaluation. An LLM relevance scorer $\mathcal{R}$ assigns an integer score $r(x) \in \{1, 2, 3, 4, 5\}$ to each paper in $\mathcal{S}_1 \cup \mathcal{S}_{\text{syn}}$ using research question query and tailored inclusion criteria which are applied independently per instance, following the same setup as (Dennstädt et al., 2024). We compute two metrics to evaluate how well an agent baseline distinguishes original papers from its indirectness-perturbed counterparts. (1) *Direct Evidence Win Rate.* For each $x_i \in \mathcal{S}_1$ and each synthetic variant x_i^c, we define a binary preference indicator $w(x_i, x_i^c) = \mathbb{1}[r(x_i) > r(x_i^c)]$, and aggregate it:

$$\text{Win Rate} = \frac{1}{|\mathcal{S}_1| \cdot |\mathcal{C}|} \sum_{x_i \in \mathcal{S}_1} \sum_{c \in \mathcal{C}} w(x_i, x_i^c) \quad (1)$$

We also measure (2) *Relevance Score Gap*, average score margin between each original paper and its indirect-evidence synthetic variants:

$$\text{Score Gap} = \frac{1}{|\mathcal{S}_1| \cdot |\mathcal{C}|} \sum_{x_i \in \mathcal{S}_1} \sum_{c \in \mathcal{C}} \big(r(x_i) - r(x_i^c)\big) \quad (2)$$

A higher score gap indicates stronger sensitivity to indirectness violations, which is a desirable property for literature review retrieval agent. Unlike Stage 1 which has relevance ground truth, Stages 2-4 are open-ended tasks with no existing ground truth. We evaluate the rest of the stages using LLM judge to asses quality dimensions described in the following sections.

3.2 Stage 2: Hypotheses Generation

Trustworthiness risk: ungrounded hypotheses. First, an LLM can generate hypotheses that are internally coherent yet lack traceable support in the literature. Second, it can appear valid only because no competing explanation was considered. Both risks are particularly dangerous in medical research, as a flawed hypothesis directly propagates

to next stage protocol design. Valid hypotheses require grounding in existing evidence (Jing et al., 2022) and a deductive checkpoint against competing explanations is then necessary before proceeding (Pietarinen and Stanley, 2025), motivating our following design.

Evidence-grounding mechanism. Given relevant literature set $\mathcal{L}$ from Stage 1, this stage operates through three sequential agents:

- *PICO Parser:* extracts structured PICO from each paper in $\mathcal{L}$, yielding a parsed corpus $\mathcal{P}$.
- *Gap Analyzer:* reads $\mathcal{P}$ to identify underexplored PICO configurations as prioritized evidence gaps $\mathcal{G}$.
- *Hypothesis Generator:* produces candidate hypotheses $\mathcal{H}$ from $\mathcal{P}$ and $\mathcal{G}$, then grounds each hypothesis h_j through two steps:
  1. *Evidence grounding:* retrieves supporting literature $\mathcal{F}_j$ by matching the PICO elements extracted from h_j, flagging any element with no coverage in $\mathcal{F}_j$.
  2. *Verification:* since evidence rarely yields a clean verdict for a single claim (Glockner et al., 2024), a competing hypothesis h_j^{comp} is generated as an alternate or confounding explanation and checked against realistic study design as in $\mathcal{F}_j$. This process yield reasoning trace r_j and a confidence score $\kappa_j \in \{\text{low, moderate, high}\}$.

Each hypotheses output h_j carrying structured evidence-grounding: supporting citations $\mathcal{F}_j$, reasoning trace r_j, and confidence score κ_j which then passed to Stage 3.

Evaluation. Each hypothesis is scored by an LLM judge (1–5 scale) on dimensions adapted from prior work (Qi et al., 2024): (1) *novelty*: the hypothesis introduces a research question not already answered by the source papers; (2) *verifiability*: scored on the full hypothesis output including competing hypotheses and confidence reasoning, capturing both PICO consistency and whether the hypothesis holds against alternative explanations. Additionally, (3) *evidence grounding*: claims must be explicitly cited from the provided literature.

3.3 Stage 3: Protocol Design

Trustworthiness risk: non-compliant protocols. An LLM can generate a protocol that appears structurally complete but violates clinical research policy standards. This is particularly dangerous as a protocol that does not meet requirements cannot proceed to institutional review regardless of its scientific quality, and the risk propagates from Stage 2 if the generator is not constrained to the hypothesis evidence chain.

Policy-compliance mechanism. Given h_j, $\mathcal{F}_j$, and r_j from Stage 2, this stage operates in two steps: *(1) protocol generator* produces a structured protocol Π_j with precise actions grounded within the PICO elements of h_j and evidence chain $\mathcal{F}_j$; *(2) compliance reviewer* evaluates Π_j against four policy frameworks (Table 2), mapping critical flags to revision targets returned to the generator. The loop terminates when no critical flags remain or a maximum iteration count is reached, yielding Π_j^* passed to Stage 4.

Evaluation. Each Π_j^* is scored by an LLM judge on dimensions adapted from Luo et al. (2025): *level of detail*, *correctness*, and additional *compliance*, each rated on a 1–5 scale.

Policy	Governed Layer
SPIRIT	**Study protocol** structure. Objectives, interventions, outcomes, randomization, sample size.
IRB	**Human subjects** protection. Consent, eligibility, risk-benefit, compensation, questionnaire design.
HIPAA *Health Insurance Portability and Accountability Act of 1996*	**Sensitive data** management. Collection, storage, de-identification, access control, sharing.
FDA SaMD *Food and Drug Administration's Software as a Medical Device*	**Software/model** deployment. Risk classification, validation, safety documentation, performance monitoring.

Table 2: Policy frameworks and corresponding layers.

3.4 Stage 4: Paper Writing

Trustworthiness risk: ungrounded, non-compliant reporting. An LLM can introduce unsupported claims or fabricate citations, and may fail to meet structured reporting standards. Both risks propagate from earlier stages if the evidence chain and compliance-reviewed protocol are not carried forward.

Evidence-grounding and policy compliance mechanism. Given Π_j^*, Stage 4 generates a six-section manuscript: background, rationale, methods, expected outcomes, limitations, and ethics. Methods section is translated directly from Π_j^* and citation leverage complete PICO-parsed papers $\mathcal{P}$. A SPIRIT Manuscript Readiness checker then evaluates reporting completeness, returning missing items to the generator until coverage is sufficient or a maximum iteration is reached.

Evaluation. Each manuscript is scored by an LLM judge on: *compliance*, *clarity*, and *evidence support* (whether claims are grounded in presented references), each rated on a 1–5 scale.

4 Experiment

4.1 Research Questions

RQ1 How does EGPC-MEDSCIENTIST improve evidence grounding and policy compliance quality?

RQ2 How does PICO indirectness and policy compliance behave across elements and frameworks?

RQ3 How consistent is LLM judge as an evaluator?

4.2 Baselines

We follow the LLM-based scientific discovery taxonomy of Zheng et al. (2025) and organize baselines into two categories: *Tool* and *AI-based scientist*. No single existing system covers all four stages with medical grounding, so we select the most suitable available baseline per stage (Table 3).

Tool baselines are purpose-built for a specific discovery stage, making them the most direct point of comparison for evaluating our stage-level improvements.

- *LLM-based SLR* (Dennstädt et al., 2024) performs title-abstract screening using tailored inclusion criteria on the same SLR dataset (De Bruin et al., 2023). It serves as the Stage 1 baseline to test whether our PICO-GRADE scorer improves indirectness detection.

- *BioHypoGen* (Qi et al., 2024) uses three sequential agents: Analyst, Scientist, and Critic. The Analyst passes a summarized background to the Scientist to generate hypotheses, and the Critic reviews them. Analyst summary do not preserve source references thus the hypotheses does not have a traceable evidence chain. It serves as the Stage 2 baseline to assess whether structured evidence-grounded hypothesis generation produces stronger grounding and verifiability.

To our knowledge, purpose-built medical tools for protocol design or manuscript generation remain underexplored (Zheng et al., 2025), so we use AI-based scientist baselines for Stage 3 and 4.

AI-based scientist baselines cover multiple stages, letting us evaluate whether existing designs carry evidence and compliance risks across stages in the medical domain.

Method	Lit. Rev.	Hypo. Gen.	Prot.	Paper	Trust. Mitigation
Biomedical discovery tool					
LLM-SLR (Dennstädt et al., 2024)	✓	-	-	-	-
BioHypoGen (Qi et al., 2024)	-	✓	-	-	-
AI-based scientist					
AI-Scientist (Lu et al., 2026)	✗	✓	-	✗	-
BioResearcher (Luo et al., 2025)	✓	✗	✓	-	-
SafeScientist (Zhu et al., 2025)	✗	✗	✗	✓	Harmful monitor, Malicious defender
EgPc-MedScientist (Ours)	✓	✓	✓	✓	PICO evidence, Clinical policy

✓ comparable stage ✗ incompatible stage - stage absent

Table 3: Coverage of research stages and trustworthiness mitigation across baselines. ✗ indicates a stage exists but produces output incompatible with our medical pipeline (e.g., AI-Scientist generates ML experiment papers, not clinical protocols)

- *BioResearcher* (Luo et al., 2025) end-to-end biomedical research framework stages: Literature Processing, Protocol Design, and iterative Experiment. It has no explicit hypothesis generation stage, and its iterative Experiment stage assumes wet or dry lab execution, which is outside the scope of our pipeline, so comparison is limited to Stage 1 and 3.

- *AI-Scientist* (Lu et al., 2026) covers Idea Generation, Experiment Execution, and Paper Write-up. Its Idea Generation iteratively evaluates candidate ideas against an existing archive and outputs an experimental design, which is the closest analog to our Stage 2. Although it is a general domain framework, its iterative self-reflection represents a strong design for considering available evidence in the idea generation loop. We test whether PICO-grounded hypothesis generation improves its evidence grounding and verifiability.

- *SafeScientist* (Zhu et al., 2025) is a safety-aware framework focused on adversarial attack defense: a Prompt Monitor filters harmful queries, a malicious-defense agent for Discussion stage, and Tool Use selection follows rather than a complete study protocol. Paper Writing applies ethical compliance checks which we use as baseline in Stage 4. In their released implementation, citations are retrieved after the paper is generated rather than built from a validated evidence chain.

Details of the implementation adaptation of each baseline on each stage are shown in Appendix B.3.

4.3 Implementation

Dataset. We use the public SLRs SYNERGY dataset (De Bruin et al., 2023) that was also used in the prior work (Dennstädt et al., 2024). The dataset contains ten previously published SLRs covering various biomedical topics. We use only the titles and abstracts of the relevant papers, totaling 1,170 papers (391 of which become our relevance ground-truth for the PRISM benchmark). The processing steps and the statistics of the datasets are shown in Appendix A.1.

Agents and LLM Judger's Model. In all stages, the MedScientist agents used Llama-3.3-70B Instruct model (Dubey et al., 2024)[1] running fully offline on a local workstation. For the LLM judger used in the evaluation, we use the Claude Sonnet 4.6[2], accessed through Anthropic's HIPAA-compliant API. All experiments are conducted on a workstation equipped with an Intel Xeon W-2235 CPU @ 3.80 GHz (12 cores) and an NVIDIA GeForce RTX A6000 GPU (48 GB). Details on model parameters and estimated costs are reported in Appendix B.2. All system prompts, both for MedScientist agents and LLM Judger for all stages, are shown in Appendix B.1.

5 Result

We report results across all four stages. RQ1 covers stage-level evidence grounding and compliance gains (§5.1). RQ2 breaks down PICO element sensitivity and policy coverage (§5.2). RQ3 examines LLM judge consistency (§5.3).

5.1 RQ1 Stage Level Evidence Grounding and Compliance Quality

Method	Relevance Score Gap ↑		Direct Evidence Win Rate ↑	
LLM-based SLR (Dennstädt et al., 2024)	0.762		0.251	
BioResearcher (Luo et al., 2025)	0.733	↓3.8%	0.294	↑17.1%
EGPC-MEDSCIENTIST	0.801	↑5.1%	**0.336**	**↑33.9%**
Ablation				
EGPC-MEDSCIENTIST w/ inclusion criteria	**0.849**	**↑11.4%**	0.275	↑9.6%

Table 4: **Stage 1 Literature review.** GRADE-PICO scorer improves indirectness detection over baselines. Colored values (%) show relative change over tool baseline. Statistical significance see App. A.3.

Stage 1. EGPC-MEDSCIENTIST improves over LLM-SLR on direct evidence win rate (+33.9%) (Table 4), showing that the GRADE-PICO scorer helps the agent better distinguish indirect-evidence papers from actual relevant ones. BioResearcher gains on win rate (+17.1%) but lags behind on score gap (-3.8%), suggesting that LLM-SLR is actually more decisive once it detects any indirectness, penalizing violations more strongly. The ablation shows that adding inclusion criteria improves decisiveness but reduces coverage. Regardless of the configuration, both variants outperform the baseline, confirming that GRADE-PICO scoring consistently improves indirectness detection in literature review stage.

Method	Novelty ↑		Evidence Grounding ↑		Verifiability ↑	
BioHypoGen (Qi et al., 2024)	1.70		2.00		2.60	
AI-Scientist (Lu et al., 2026)	**3.60**	↑1.90	2.10	↑0.10	4.10	↑1.50
EGPC-MEDSCIENTIST	3.00	↑1.30	**3.20**	↑1.20	**4.40**	↑1.80
Ablation						
EGPC-MEDSCIENTIST w/ verifier	2.80	↑1.10	3.00	↑1.00	4.10	↑1.50

Table 5: **Stage 2 Hypothesis generation.** EGPC-MEDSCIENTIST achieves the highest evidence grounding and verifiability across all methods. Colored values (5-point scale) show absolute improvement compared to the tool baseline. Statistical significance see App. A.3.

Stage 2. BioHypoGen scores consistently low across all dimensions (Table 5). The LLM judge consistently found that Hypotheses generated by BioHypoGen merely summarize knowledge already covered in the source papers with broad narrative statements. AI-Scientist achieves the highest novelty (3.60 vs. 3.00), consistent with its iterative self-reflection, but EGPC-MEDSCIENTIST outperforms on evidence grounding (3.20 vs. 2.10) and verifiability (4.40 vs. 4.10). In clinical research, a novel but ungrounded or unverifiable hypothesis does not advance scientific practice (Jing et al., 2022; Pietarinen and Stanley, 2025), making these two dimensions more critical than novelty at this stage. The ablation confirms our design on the verification step (§3.2), which improves both evidence grounding (3.00 → 3.20) and verifiability (4.10 → 4.40), showing that grounding competing hypotheses against the same evidence strengthens not only verifiability but overall evidence quality.

Stage 3. EGPC-MEDSCIENTIST outperforms BioResearcher across all three dimensions: details (+0.50), correctness (+0.30), and compliance

[1] https://hf.co/meta-llama/Llama-3.3-70B-Instruct
[2] https://www.anthropic.com/news/claude-sonnet-4-6

Method	Details ↑	Correctness ↑	Compliance ↑
BioResearcher (Luo et al., 2025)	3.45	3.20	3.20
EGPC-MEDSCIENTIST	**3.95**	**3.50**	**4.10**

Table 6: **Stage 3 Protocol design** evaluation across details, correctness, and compliance. Statistical significance see App. A.3.

(+0.90), with compliance showing the largest gain (Table 6). The strong compliance gain confirms that an explicit policy review loop is needed; without it, protocols may appear structurally sound while still violating clinical research requirements such as patient data handling (HIPAA) or human subjects protection (IRB).

Method	Clarity ↑	Evidence support ↑	Compliance ↑
SafeScientist (Zhu et al., 2025)	3.60	2.33	2.50
EGPC-MEDSCIENTIST	**3.90**	**3.20**	**3.30**

Table 7: **Stage 4 Paper writing** evaluation across three quality dimensions: clarity, evidence support, and compliance. Statistical significance see App. A.3.

Stage 4. EGPC-MEDSCIENTIST outperforms SafeScientist across all three dimensions: compliance (+0.80), clarity (+0.30), and evidence support (+0.87) as seen in Table 7. The compliance and evidence support gains are the most notable, and are consistent with SafeScientist's back-to-front design (§4.2), where citations are retrieved after the paper is generated rather than built from a recorded evidence chain. This confirms that carrying forward supporting citations $\mathcal{P}$ (§3.4), whose papers are validated through PICO screening and hypothesis matching in earlier stages, allowed traceability, which improves evidence support quality. This is further illustrated by SafeScientist producing zero cited sentences on two topics despite Semantic Scholar access (Figure 5b, §5.2)

5.2 RQ2 PICO Sensitivity, Compliance Coverage, and Citation Support

PICO sensitivity. Figure 3 shows the impact of our GRADE-PICO scorer on each baseline across both metrics. Adding the scorer consistently improves indirectness detection across all PICO elements, with the largest gain on ***comparison*** for BioResearcher (0.17 → 0.43), the hardest element to detect across all methods (Guyatt et al., 2011). A tradeoff appears where BioResearcher's ***outcome*** sensitivity drops (0.61 → 0.37) as attention is redistributed across all four elements. On score gap, LLM-SLR is more decisive when it catches a violation while BioResearcher catches more violations but penalizes them less strongly. EGPC-MEDSCIENTIST gives the best combination of both higher coverage and clear penalization. Per-dataset details available in Appendix C Figure 32.

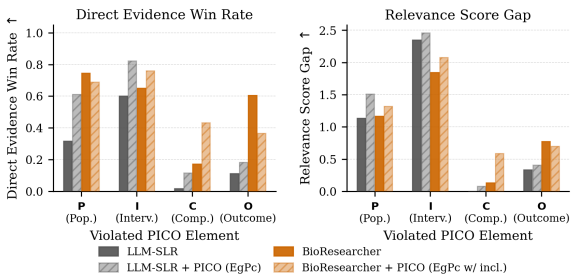

Figure 3: **GRADE-PICO scorer consistently improves all baselines across PICO elements**, with largest gain (0.18 → 0.43) even on the hardest element *Comparison* (Guyatt et al., 2011). Higher is better (↑).

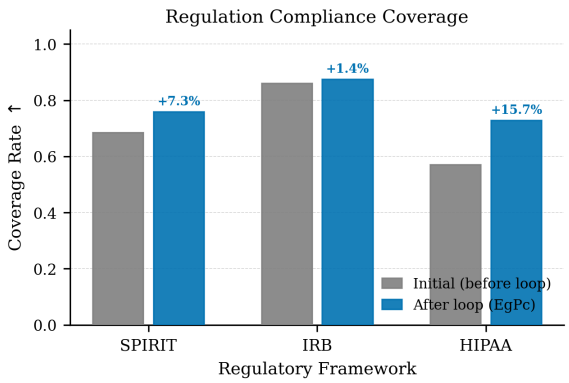

Figure 4: Policy compliance coverage **improvement after EGPC-MEDSCIENTIST review loop**, averaged across 10 medical topics.

Compliance coverage. Our compliance review loop consistently improves coverage across policy frameworks (Figure 4), with the largest gain on HIPAA (+15.7%), indicating that data privacy requirements are mostly missed without explicit policy grounding. IRB starts already high (0.861) and gains minimally, as clinical trial protocols naturally address human subjects requirements even without the loop. FDA SaMD is excluded as it applies to <20% of protocols, though it becomes relevant for topics involving software or model deployment such as PTSD symptom modeling (van de Schoot et al., 2018) and cardiovascular risk prediction (Wolters et al., 2018).

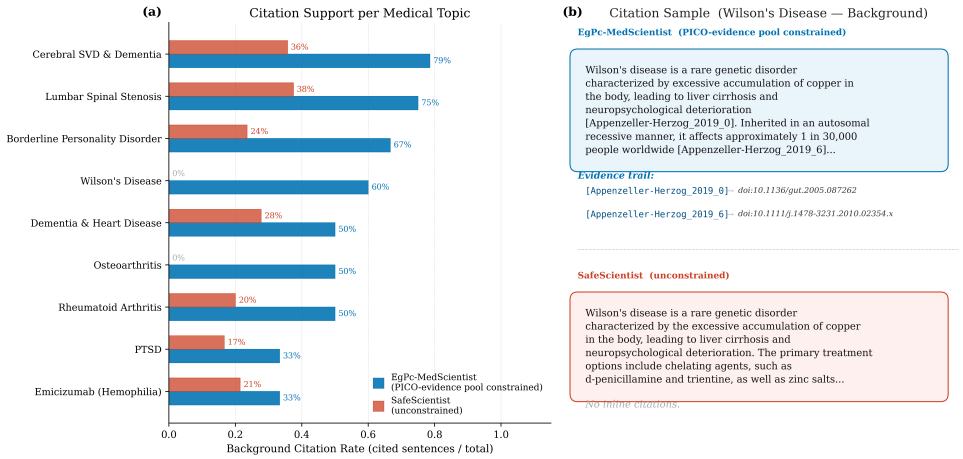

Figure 5: (a) Final manuscript shows **EGPC-MEDSCIENTIST provide more citation support** across all medical topics, compared to SafeScientist for generated statements in background section. In Wilson's Disease topic (b), **EGPC-MEDSCIENTIST provide verifiable sources** while SafeScientist failed to provide any citation

Citation support. EGPC-MEDSCIENTIST provides higher citation support across all topics (Figure 5), where more statements are backed by traceable sources (Gao et al., 2023). SafeScientist produces 0% cited sentences on Wilson's Disease and Osteoarthritis despite having Semantic Scholar access, suggesting tool access alone does not guarantee evidence-supported reporting.

5.3 RQ3 LLM-Judger Consistency, Human Validation.

LLM-Judger Consistency. To validate our evaluation methodology, we provide comparison result of our primary judge (Claude Sonnet 4.6) against GPT(5.4)-judger across baselines and metrics of Stage 2. Figure 6 shows consistent ranking across baselines, with EGPC-MEDSCIENTIST highest on Evidence Grounding and Verifiability, AI-Scientist on Novelty, and BioHypoGen consistently last.

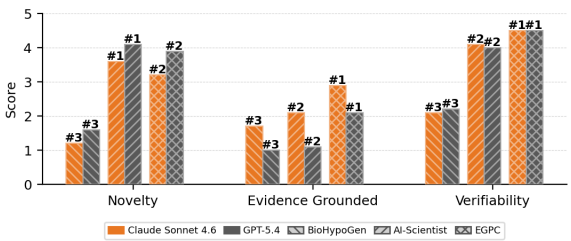

Figure 6: **Hypothesis Generation ratings between two LLM judgers**, methods ordering remain consistent.

Human validation on hypothesis quality. To evaluate whether structured evidence output leads to more consistent quality judgments, we conducted a human evaluation with PhD students in medical education and public health. The full structured output (citations, competing hypotheses, confidence reasoning) increased inter-annotator agreement from Krippendorff's α <0.3 to >0.67, compared to hypothesis statements alone. In a follow-up discussion, annotators pointed out some problems, such as hypothesis combining multiple populations into one statement, reading as a paper fragment rather than a testable claim, and all consistently matched baseline without PICO grounding, while annotators remained blind to method identity throughout. This indicates that evidence grounding helps not only automated evaluation but also human judgment of hypothesis quality (Appendix A.2).

6 Conclusion

We presented EGPC-MEDSCIENTIST, a four-stage medical research automation framework with explicit trustworthiness mechanisms targeting evidence grounding and policy compliance, along with PRISM, a clinician-validated benchmark for indirect evidence detection in literature review. Experiments show measurable trustworthiness gaps in existing pipelines across all stages, and that EGPC-MEDSCIENTIST improves evidence grounding, verifiability, and compliance over baselines. A recurring pattern behind these gaps is the generate-first, cite-later design assumption, which our Stage 4 results show has measurable costs in evidence grounding and compliance. We hope this motivates the community to treat domain-specific trustworthiness as a first requirement when adapting research automation to medical and clinical settings.

Limitations

Benchmark scope. PRISM is constructed from title and abstract pairs only, following the screening setting of Dennstädt et al. (2024). Full-text papers contain richer PICO information, and indirectness violations that are detectable from an abstract may be harder to spot or easier to miss when the full text provides more context. Future work should examine whether GRADE-PICO scoring remains effective under full-text conditions.

Synthetic data generation. The counterfactual instances in PRISM are generated by an LLM (Claude Sonnet 4.6) and are thus constrained by what the model can produce as a realistic clinical concept. While the clinician verification confirms a 96.9% violation rate, the synthetic variants may not cover the full range of indirectness patterns that appear in real clinical literature, particularly for subtle population or comparator mismatches.

Pipeline scope. EGPC-MEDSCIENTIST covers the research design pipeline up to protocol paper writing. It does not execute experiments, collect data, or analyze results. This means the compliance and evidence grounding properties we measure apply to the planning stages only.

Stage-level evaluation. Each stage is evaluated independently using its own baseline and metrics. We do not evaluate the full pipeline end-to-end, meaning we do not measure how errors or quality gaps from one stage propagate and compound across subsequent stages. For example, a borderline hypothesis from Stage 2 may still produce a structurally complete protocol in Stage 3, but whether the overall research output remains trustworthy under such conditions is not captured by our per-stage metrics. End-to-end evaluation with a shared quality measure across all four stages is left for future work.

Medical topic coverage. The 10 topics in the SYNERGY dataset cover a range of biomedical areas but are not representative of all clinical domains. Highly specialized fields may require additional or different PICO configurations and policy considerations that are not captured in our current benchmark.

Ethical Statement

Intended use and limitations of safeguards. EGPC-MEDSCIENTIST is a research assistance tool, not a replacement for clinical expertise. While the PICO grounding and compliance review mechanisms reduce the risk of ungrounded evidence and policy violations, they do not eliminate it. Outputs at any stage, particularly protocol designs and manuscript drafts, must be reviewed by qualified researchers and clinicians before real-world use. Any deployment should include mandatory expert review at each stage.

References

ACL. 2026. Acl statement on desk rejecting papers with hallucinated references. ACL Website. Accessed: May 2026.

Manar Aljohani, Jun Hou, Sindhura Kommu, and Xuan Wang. 2025. A comprehensive survey on the trustworthiness of large language models in healthcare. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 6720–6748, Suzhou, China. Association for Computational Linguistics.

Christian Appenzeller-Herzog, Tim Mathes, Marlies L S Heeres, Karl Heinz Weiss, Roderick H. J. Houwen, and Hannah Ewald. 2019. Comparative effectiveness of common therapies for wilson disease: A systematic review and meta-analysis of controlled studies. *Liver International*, 39:2136 – 2152.

Nokimul Hasan Arif, Qian Lou, and Mengxin Zheng. 2026. Conjunctive prompt attacks in multi-agent LLM systems. *CoRR*, abs/2604.16543.

Yoshua Bengio, Michael K. Cohen, Damiano Fornasiere, Joumana Ghosn, Pietro Greiner, Matt MacDermott, Sören Mindermann, Adam Oberman, Jesse Richardson, Oliver Ethan Richardson, Marc-Antoine Rondeau, Pierre-Luc St-Charles, and David Williams-King. 2025. Superintelligent agents pose catastrophic risks: Can scientist ai offer a safer path? *ArXiv*, abs/2502.15657.

Daniel Bos, Frank J. Wolters, Sirwan K. L. Darweesh, Meike W. Vernooij, Frank de Wolf, M. Arfan Ikram, and Albert Hofman. 2018. Cerebral small vessel disease and the risk of dementia: A systematic review and meta-analysis of population-based evidence. *Alzheimer's & Dementia*, 14.

Jonathan De Bruin, Yongchao Ma, Gerbrich Ferdinands, Jelle Teijema, and Rens Van de Schoot. 2023. SYNERGY - Open machine learning dataset on study selection in systematic reviews.

Fabio Dennstädt, Johannes Zink, Paul Martin Putora, Janna Hastings, and Nikola Cihoric. 2024. Title and abstract screening for literature reviews using large language models: an exploratory study in the biomedical domain. *Systematic reviews*, 13(1):158.

Anouk A. M. T. Donners, Carin M. A. Rademaker, Lisanne A. H. Bevers, Alwin D. R. Huitema, Roger

E. G. Schutgens, Toine C. G. Egberts, and Kathelijn Fischer. 2021. Pharmacokinetics and associated efficacy of emicizumab in humans: A systematic review. *Clinical Pharmacokinetics*, 60:1395 – 1406.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.

Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling large language models to generate text with citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.

Max Glockner, Ieva Staliūnaitė, James Thorne, Gisela Vallejo, Andreas Vlachos, and Iryna Gurevych. 2024. AmbiFC: Fact-checking ambiguous claims with evidence. *Transactions of the Association for Computational Linguistics*, 12:1–18.

Gordon Henry Guyatt, Andrew David Oxman, Regina Kunz, James Woodcock, Jan L. Brożek, Mark Helfand, Pablo Alonso-Coello, Yngve Falck-Ytter, Roman Jaeschke, Gunn Elisabeth Vist, Elie A. Akl, Piet N. Post, Susan L. Norris, Joerg J. Meerpohl, Vijay K Shukla, Mona Nasser, and Holger J. Schunemann. 2011. Grade guidelines: 8. rating the quality of evidence–indirectness. *Journal of clinical epidemiology*, 64 12:1303–10.

Rochana Prih Hastuti, Rian Adam Rajagede, Mansour Al Ghanim, Mengxin Zheng, and Qian Lou. 2025. Factuality beyond coherence: Evaluating LLM watermarking methods for medical texts. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, Suzhou, China. Association for Computational Linguistics.

Wenyue Hua, Xianjun Yang, Mingyu Jin, Zelong Li, Wei Cheng, Ruixiang Tang, and Yongfeng Zhang. 2024. TrustAgent: Towards safe and trustworthy LLM-based agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10000–10016, Miami, Florida, USA. Association for Computational Linguistics.

Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, Hanchi Sun, Zhengliang Liu, Yixin Liu, Yijue Wang, Zhikun Zhang, Bertie Vidgen, Bhavya Kailkhura, Caiming Xiong, and 52 others. 2024. Position: TrustLLM: Trustworthiness in large language models. In *Forty-first International Conference on Machine Learning*.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabsa. 2023. Llama guard: Llm-based input-output safeguard for human-ai conversations. *CoRR*, abs/2312.06674.

Madhan Jeyaraman, Sathish Muthu, and Parvez Ahmad Ganie. 2020. Does the source of mesenchymal stem cell have an effect in the management of osteoarthritis of the knee? meta-analysis of randomized controlled trials. *CARTILAGE*, 13:1532S – 1547S.

Xia Jing, Vimla L Patel, James J Cimino, Jay H Shubrook, Yuchun Zhou, Chang Liu, and Sonsoles De Lacalle. 2022. The roles of a secondary data analytics tool and experience in scientific hypothesis generation in clinical research: Protocol for a mixed methods study. *JMIR Res Protoc*, 11(7):e39414.

Jitendra Jonnagaddala and Zoie Shui-Yee Wong. 2025. Privacy preserving strategies for electronic health records in the era of large language models. *npj Digit. Medicine*, 8(1).

Sebastian Joseph, Lily Chen, Jan Trienes, Hannah Göke, Monika Coers, Wei Xu, Byron Wallace, and Junyi Jessy Li. 2024. FactPICO: Factuality evaluation for plain language summarization of medical evidence. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8437–8464, Bangkok, Thailand. Association for Computational Linguistics.

Cathalijn H. C. Leenaars, Frans R. Stafleu, David De Jong, M F van Berlo, Tijmen Geurts, Tineke Coenen de Roo, Jan-Bas Prins, Rosalie Wm Kempkes, Janneke Elzinga, André Bleich, Rob B.M. de Vries, Franck L. B. Meijboom, and Merel Ritskes-Hoitinga. 2020. A systematic review comparing experimental design of animal and human methotrexate efficacy studies for rheumatoid arthritis: Lessons for the translational value of animal studies. *Animals : an Open Access Journal from MDPI*, 10.

Binglan Li, Anil Kumar Saini, Jose Guadalupe Hernandez, and Jason H. Moore. 2026. Agentic ai and the rise of in silico team science in biomedical research. *Nature biotechnology*.

Aiwei Liu, Leyi Pan, Yijian Lu, Jingjing Li, Xuming Hu, Xi Zhang, Lijie Wen, Irwin King, Hui Xiong, and Philip S. Yu. 2025. A survey of text watermarking in the era of large language models. *ACM Comput. Surv.*, 57(2):47:1–47:36.

Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2026. The ai scientist: Towards fully automated open-ended scientific discovery. *Nature*, 651:914–919.

Yi Luo, Linghang Shi, Yihao Li, Aobo Zhuang, Yeyun Gong, Ling Liu, and Chen Lin. 2025. From intention to implementation: automating biomedical research via llms. *Science China Information Sciences*, 68(7):170105.

Rosanne W. Meijboom, Helga Gardarsdottir, Toine C. G. Egberts, and Thijs J. Giezen. 2021. Patients retransitioning from biosimilar tnfα inhibitor to the corresponding originator after initial transitioning to the biosimilar: A systematic review. *Biodrugs*, 36:27 – 39.

Kaijie Mo, Siddhartha Venkatayogi, Chantal Shaib, Ramez Kouzy, Wei Xu, Byron C. Wallace, and Junyi Jessy Li. 2026. Faithfulness vs. safety: Evaluating LLM behavior under counterfactual medical evidence. *CoRR*, abs/2601.11886.

Matthijs Oud, Arnoud Arntz, Marleen L.M. Hermens, Rogier E. J. Verhoef, and Tim Kendall. 2018. Specialized psychotherapies for adults with borderline personality disorder: A systematic review and meta-analysis. *The Australian and New Zealand Journal of Psychiatry*, 52:949 – 961.

Ahti-Veikko Pietarinen and Donald E Stanley. 2025. The logic of medical reasoning: toward an integrated inductive, deductive, and abductive approach to clinical practices. *Philosophy, Ethics, and Humanities in Medicine*, 20(1):16.

Biqing Qi, Kaiyan Zhang, Kai Tian, Haoxiang Li, Zhang-Ren Chen, Sihang Zeng, Ermo Hua, Hu Jinfang, and Bowen Zhou. 2024. Large language models as biomedical hypothesis generators: A comprehensive evaluation. In *First Conference on Language Modeling*.

Muthu Sathish and Eswar Ramakrishnan. 2021. Fragility analysis of statistically significant outcomes of randomised control trials in spine surgery: A systematic review. *Spine*.

Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu, and Emad Barsoum. 2025. Agent laboratory: Using LLM agents as research assistants. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5977–6043, Suzhou, China. Association for Computational Linguistics.

Holger J. Schünemann, Ignacio Neumann, Sue Brennan, Joerg Meerpohl, Miranda Langendam, Marina Davoli, Pablo Alonso Coello, Philipp Dahm, Jun Xia, Nicole Skoetz, and Elie Akl. 2024. Indirectness. In *The GRADE Book*, version 1.0 (updated september 2024) edition. The GRADE Working Group. Last modified: 23 Oct 2025.

Yu-Zhe Shi, Mingchen Liu, Fanxu Meng, Qiao Xu, Zhangqian Bi, Kun He, Lecheng Ruan, and Qining Wang. 2025. Hierarchically encapsulated representation for protocol design in self-driving labs. In *The Thirteenth International Conference on Learning Representations*.

Xiangru Tang, Qiao Jin, Kunlun Zhu, Tongxin Yuan, Yichi Zhang, Wangchunshu Zhou, Meng Qu, Yilun Zhao, Jian Tang, Zhuosheng Zhang, and 1 others. 2025. Risks of ai scientists: prioritizing safeguarding over autonomy. *Nature Communications*, 16(1):8317.

Qwen Team. 2024. Qwen2.5: A party of foundation models.

Rens van de Schoot, Marit Sijbrandij, Sarah Depaoli, Sonja D. Winter, Miranda Olff, and Nancy E Van Loey. 2018. Bayesian ptsd-trajectory analysis with informed priors based on a systematic literature search and expert elicitation. *Multivariate Behavioral Research*, 53:267 – 291.

Juraj Vladika, Mahdi Dhaini, and Florian Matthes. 2025. Facts fade fast: Evaluating memorization of outdated medical knowledge in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 9161–9174, Suzhou, China. Association for Computational Linguistics.

Frank J. Wolters, Reffat A. Segufa, Sirwan K. L. Darweesh, Daniel Bos, M Arfan Ikram, Behnam Sabayan, Albert Hofman, and Sanaz Sedaghat. 2018. Coronary heart disease, heart failure, and the risk of dementia: A systematic review and meta-analysis. *Alzheimer's & Dementia*, 14.

Jiaqi Xue, Qian Lou, and Mengxin Zheng. 2024. BadFair: Backdoored fairness attacks with group-conditioned triggers. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8257–8270, Miami, Florida, USA. Association for Computational Linguistics.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 40 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Haofei Yu, Keyang Xuan, Fenghai Li, Kunlun Zhu, Zijie Lei, Jiaxun Zhang, Ziheng Qi, Kyle Richardson, and Jiaxuan You. 2025a. TinyScientist: An interactive, extensible, and controllable framework for building research agents. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 558–590, Suzhou, China. Association for Computational Linguistics.

Miao Yu, Fanci Meng, Xinyun Zhou, Shilong Wang, Junyuan Mao, Linsey Pan, Tianlong Chen, Kun Wang, Xinfeng Li, Yongfeng Zhang, Bo An, and Qingsong Wen. 2025b. A survey on trustworthy llm agents: Threats and countermeasures. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, KDD '25, page 6216–6226, New York, NY, USA. Association for Computing Machinery.

Hang Zhang, Qian Lou, and Yanshan Wang. 2025. Towards safe ai clinicians: A comprehensive study on large language model jailbreaking in healthcare. *arXiv preprint arXiv:2501.18632*.

Yu Zhang, Xiusi Chen, Bowen Jin, Sheng Wang, Shuiwang Ji, Wei Wang, and Jiawei Han. 2024. A comprehensive survey of scientific large language models

and their applications in scientific discovery. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8783–8817, Miami, Florida, USA. Association for Computational Linguistics.

Tianshi Zheng, Zheye Deng, Hong Ting Tsang, Weiqi Wang, Jiaxin Bai, Zihao Wang, and Yangqiu Song. 2025. From automation to autonomy: A survey on large language models in scientific discovery. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 17733–17750, Suzhou, China. Association for Computational Linguistics.

Yujun Zhou, Jingdong Yang, Yue Huang, Kehan Guo, Zoe Emory, Bikram Ghosh, Amita Bedar, Sujay Shekar, Zhenwen Liang, Pin-Yu Chen, and 1 others. 2026. Benchmarking large language models on safety risks in scientific laboratories. *Nature Machine Intelligence*, pages 1–12.

Kunlun Zhu, Jiaxun Zhang, Ziheng Qi, Nuoxing Shang, Zijia Liu, Peixuan Han, Yue Su, Haofei Yu, and Jiaxuan You. 2025. SafeScientist: Enhancing AI scientist safety for risk-aware scientific discovery. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2317, Suzhou, China. Association for Computational Linguistics.

A Method Details

A.1 PRISM PICO counterfactual benchmark construction.

We use the SYNERGY dataset (De Bruin et al., 2023)[3] to create PRISM benchmark for evaluating the relevance scoring under indirectness violations in medical SLR. We take 10 datasets across different medical topics, following (Dennstädt et al., 2024). Each dataset contains paper metadata, including title, abstract, and doi. We clean up the dataset by removing entries with an empty title or abstract. The data distribution after cleaning is shown on the left side of Table 8. In total, there are 28,036 papers, of which 1,170 (4.2%) are labeled as relevant to their respective topics. We only use the variable title, abstract, and doi.

In the next step, we use Claude Sonnet 4.6[4] to check the completeness of PICO elements in the title and abstract of each relevant paper (Label 1). We discard papers that Claude cannot detect all PICO elements in the title and abstract. The number of relevant papers for each dataset after this step is shown in Table 8, column PRISM/True Label 1, with a total of 391 papers. From these 391 papers, we generate synthetic data by modifying each PICO element in the paper, making it indirect. We use Claude Sonnet for this task using the prompt shown in Figure 7 and Figure 8, resulting in 1,564 new data points shown in Table 8, column PRISM/True Label 0.

A.2 Human Evaluation

PRISM verification is conducted by a clinician expert, assessing synthetic instances to correctly reflect the intended PICO violations (sample shown in Figure 9). We sampled 80 instances using stratified sampling covering 10 datasets and 4 PICO components (64 synthetic violated instances and 16 orignial instances as control, data distribution see Table 8). We report the Violation confirmation rate which is the proportion of synthetic instances confirmed as violated. Expert violation confirmation rate = 96.9%, with one control instance was incorrectly flagged as a violation (likely due to the complexity of the abstract resulting a reading fatigue or attention slip).

Hypothesis generation quality. Two PhD students (medical education and public health) and

[3] https://github.com/asreview/synergy-dataset
[4] https://www.anthropic.com/news/claude-sonnet-4-6

Dataset	Total	Label 0	Label 1	PRISM	
				True Label 0 (synthetic)	True Label 1 (original)
Wilson's Disease (Appenzeller-Herzog et al., 2019)	2192	2167 (98.86%)	25 (1.14%)	60	15
Cerebral small vessel disease & dementia (Bos et al., 2018)	4816	4806 (99.79%)	10 (0.21%)	12	3
Emicizumab (Donners et al., 2021)	249	235 (94.38%)	14 (5.62%)	24	6
Osteoarthritis (Jeyaraman et al., 2020)	1171	1075 (91.80%)	96 (8.20%)	372	93
Rheumatoid arthritis (Leenaars et al., 2020)	6793	6217 (91.52%)	576 (8.48%)	400	100
TNFα-inhibitors and biosimilars (Meijboom et al., 2021)	825	789 (95.64%)	36 (4.36%)	140	35
Spine surgery (Sathish and Ramakrishnan, 2021)	2695	2359 (87.53%)	336 (12.47%)	400	100
Borderline personality disorder (Oud et al., 2018)	935	915 (97.86%)	20 (2.14%)	80	20
PTSD (van de Schoot et al., 2018)	4467	4429 (99.15%)	38 (0.85%)	36	9
Dementia and heart disease (Wolters et al., 2018)	3893	3874 (99.51%)	19 (0.49%)	40	10
Total	**28036**	**26866 (95.8%)**	**1170 (4.2%)**	**1564 (80%)**	**391 (20%)**

Table 8: Distribution of the SYNERGY Dataset and PRISM.

You are an expert in evidence-based medicine and the GRADE framework for assessing certainty of evidence. Your task is to introduce "indirectness" into a clinical study abstract by modifying ONE specific PICO element. Indirectness in GRADE means the evidence does not directly apply to the review question because one of the PICO elements differs from what was intended in the review.

IMPORTANT RULES:

1. Only modify the ONE PICO element specified – keep everything else as close to the original as possible.
2. The modification must be clear and specific – a medical expert should immediately recognize it as wrong for the original study context. Do NOT make vague, borderline, or ambiguous changes.
3. The modified element must still be a real, specific clinical concept – just clearly mismatched with the original study context.
4. Preserve the writing style, structure, and length of the original abstract.
5. Do NOT add commentary or notes – output ONLY the modified title and abstract.
6. The title should also be updated to reflect the PICO modification you made.

Below is the title and abstract of a clinical study. Introduce indirectness by modifying the `{pico}` element. `{pico_instruction}`

The change must be clear and unambiguous. A medical expert reading the modified abstract should immediately recognize that the `{pico}` is wrong for the original study – not subtle, not borderline, but clearly and concretely incorrect. Choose a replacement that is a real clinical concept but obviously mismatched with the rest of the abstract.

Title: `{title}`
Original abstract: `{abstract}`

Output the modified title, modified abstract, and a one-sentence description of what you changed.

Figure 7: The prompt used for generating indirectness. The variable `pico` represents one of the four components ("Population", "Intervention", "Comparator", or "Outcome"), and `pico_instruction` provides the specific modification instruction for each element as shown in Figure 8.

"P": "Population – change the patient population to a clearly inappropriate or mismatched group that would make the study unreliable for the original research question (e.g., knee OA patients -> healthy volunteers, adult cartilage repair -> pediatric bone fracture patients, osteoarthritis patients -> cancer patients)",

"I": "Intervention – change the intervention to a clearly different and inappropriate treatment that does not address the original condition (e.g., stem cell therapy -> dietary supplements, chondrocyte implantation -> physiotherapy alone, intra-articular injection -> oral painkillers)",

"C": "Comparator – change the comparator to something inappropriate or misleading that makes the comparison invalid (e.g., placebo-controlled -> compared against an unrelated surgical procedure, matched control -> no control group mentioned, sham procedure -> historical data from a different condition)",

"O": "Outcome – change the primary outcome to something irrelevant or untrustworthy that does not measure what the study claims to measure (e.g., cartilage regeneration -> patient satisfaction survey, functional knee score -> blood biomarker levels, MRI-confirmed healing -> self-reported mood improvement)",

Figure 8: PICO-specific modification instruction

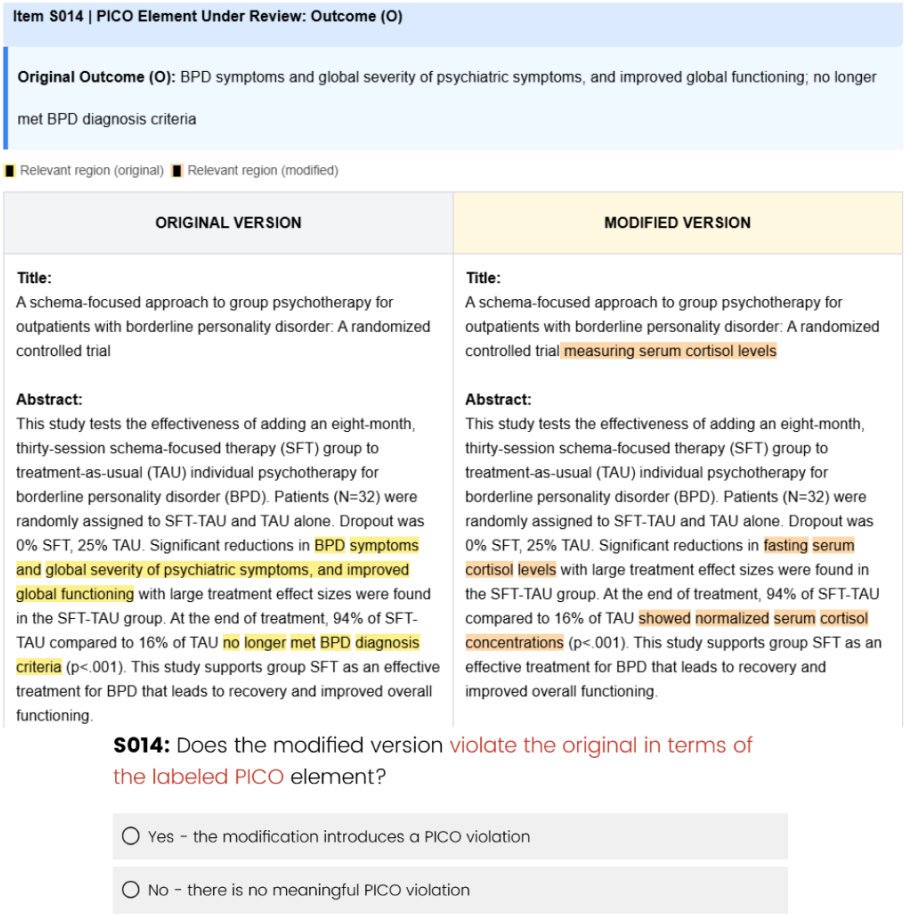

Item S014 | PICO Element Under Review: Outcome (O)

Original Outcome (O): BPD symptoms and global severity of psychiatric symptoms, and improved global functioning; no longer met BPD diagnosis criteria

Relevant region (original) Relevant region (modified)

ORIGINAL VERSION	MODIFIED VERSION
Title: A schema-focused approach to group psychotherapy for outpatients with borderline personality disorder: A randomized controlled trial	**Title:** A schema-focused approach to group psychotherapy for outpatients with borderline personality disorder: A randomized controlled trial measuring serum cortisol levels
Abstract: This study tests the effectiveness of adding an eight-month, thirty-session schema-focused therapy (SFT) group to treatment-as-usual (TAU) individual psychotherapy for borderline personality disorder (BPD). Patients (N=32) were randomly assigned to SFT-TAU and TAU alone. Dropout was 0% SFT, 25% TAU. Significant reductions in BPD symptoms and global severity of psychiatric symptoms, and improved global functioning with large treatment effect sizes were found in the SFT-TAU group. At the end of treatment, 94% of SFT-TAU compared to 16% of TAU no longer met BPD diagnosis criteria (p<.001). This study supports group SFT as an effective treatment for BPD that leads to recovery and improved overall functioning.	**Abstract:** This study tests the effectiveness of adding an eight-month, thirty-session schema-focused therapy (SFT) group to treatment-as-usual (TAU) individual psychotherapy for borderline personality disorder (BPD). Patients (N=32) were randomly assigned to SFT-TAU and TAU alone. Dropout was 0% SFT, 25% TAU. Significant reductions in fasting serum cortisol levels with large treatment effect sizes were found in the SFT-TAU group. At the end of treatment, 94% of SFT-TAU compared to 16% of TAU showed normalized serum cortisol concentrations (p<.001). This study supports group SFT as an effective treatment for BPD that leads to recovery and improved overall functioning.

S014: Does the modified version violate the original in terms of the labeled PICO element?

- ○ Yes - the modification introduces a PICO violation
- ○ No - there is no meaningful PICO violation

Figure 9: PRISM synthetic dataset human verification with clinician expert.

one clinician rated all 20 hypothesis outputs (10 topics × 2 methods) under two conditions. Task 1 (Figure10) showed only the hypothesis statement, asking whether it was specific, clinically meaningful, and testable. Task 2 showed the full output including evidence chain and competing hypotheses, asking whether the hypothesis represents an open and meaningful research gap (Figure11 and Figure12). EGPC-MEDSCIENTIST-generated hypotheses rated as more verifiable research gaps compared to baseline (mean = 2.80 vs 1.40, scale 1–3).

Respondents were recruited through academic social network. All participation was voluntary and was explicitly stated in the consent page of the questionnaire.

A.3 Statistical Significance

To verify result consistency across the 10 medical topics, we examine the standard deviation of evaluation scores. For Stages 1, 3, and 4, EGPC-MEDSCIENTIST shows equal or lower standard deviation compared to all baselines, indicating stable improvements across diverse medical domains. For Stage 2, EGPC-MEDSCIENTIST shows standard deviation around 0.6 across metrics, which remains modest relative to the 1–5 scoring scale (less than 15% of the full scale range). This variation is expected, as higher-performing methods naturally exhibit greater spread.

Statistical significance was further assessed using one-sided paired Wilcoxon signed-rank tests, comparing our proposed methods against baseline methods across all four stages. As shown in Table 9, our proposed methods are statistically significant on key evaluation metrics at each stage ($p < 0.05$). The significant differences hold across diverse baseline methods and evaluation dimensions, supporting that the improvements are reliable and not due to random variation across the tested datasets.

Stage	Baseline	Metric	p-value
Lit. Review	LLM-based SLR	Win Rate	0.043*
	BioResearcher	Win Rate	0.038*
Hypothesis Gen.	BioHypoGen	Novelty	0.001**
		Evidence Ground.	0.005**
		Verifiability	0.001**
	AI-Scientist	Evidence Ground.	0.011*
Protocol Design	BioResearcher	Details	0.037*
		Compliance	0.003**
Paper Writing	SafeScientist	Completeness	0.012*
		Clarity	0.042*
		Compliance	0.015*

$^{*}p < 0.05$, $^{**}p < 0.01$

Table 9: Statistically significant comparisons between proposed methods and baselines across all stages (one-sided paired Wilcoxon signed-rank test). Only significant results ($p < 0.05$) are shown.

B Implementation

B.1 Prompts System

Stage 1: PICO-aware instruction. In Figure 13, we can see the PICO-aware instruction that later will be augmented to the base prompt. In addition to the instruction shown, we also ask LLM to explicitly write each PICO element in the JSON output.

Stage 2: Hypotheses Generation. Figure 14 and Figure 15 show the PICO parser and the gap analyzer agent, respectively. Figure 16 and Figure 17 show the hypotheses generator and verifier, respectively.

Stage 3: Compliance reviewer. Figure 18 and Figure 19 show the prompt that we use at compliance review and revise in Stage 3. For the protocol generator, we use the same prompt as BioResearcher (Luo et al., 2025).

Stage 4: Manuscript reviewer. Figure 23 to Figure 28 are prompts used for generating each section of the protocol paper. While Figure 29 is used for refining the whole manuscript. The concept of creating agents for each section is inspired by (Zhu et al., 2025; Yu et al., 2025a). After the paper is generated, we use the prompts in Figures 30 and Figure 31 to review the compliance coverage.

LLM Judger prompts used in the evaluation for each stage can be seen in the Figure 20, Figure 21, and Figure 22. In addition to the prompts shown, we also include a valid JSON example at the end.

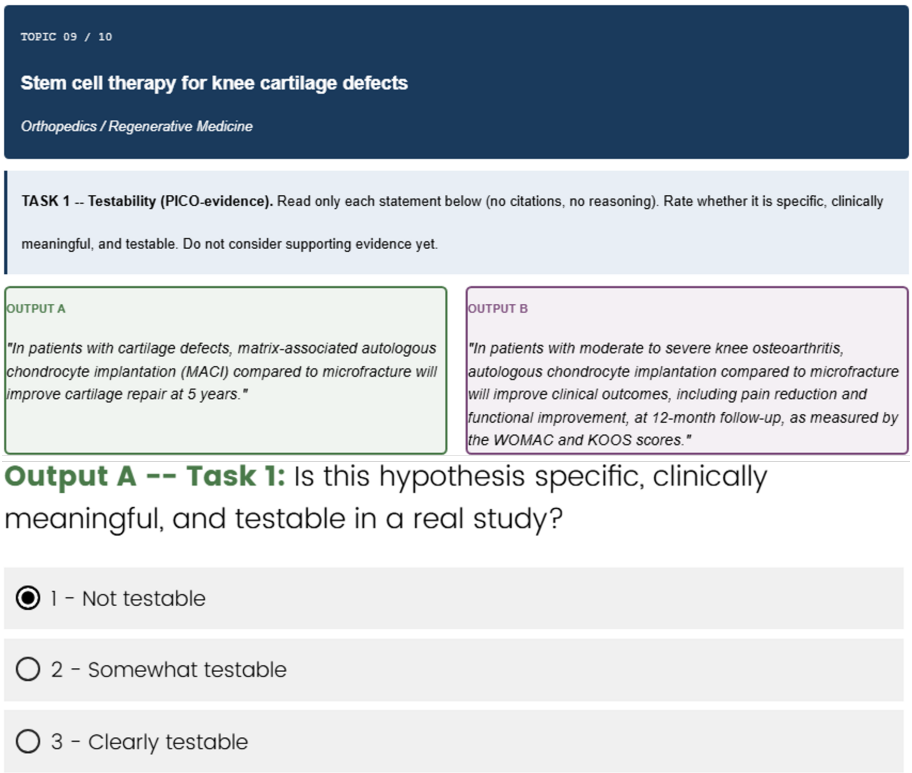

Figure 10: **Hypothesis generation human evaluation. Task 1: Testability**, It asked whether the hypothesis by itself was specific, clinically meaningful, and testable according to PICO (Population Intervention Comparison and Outcome).

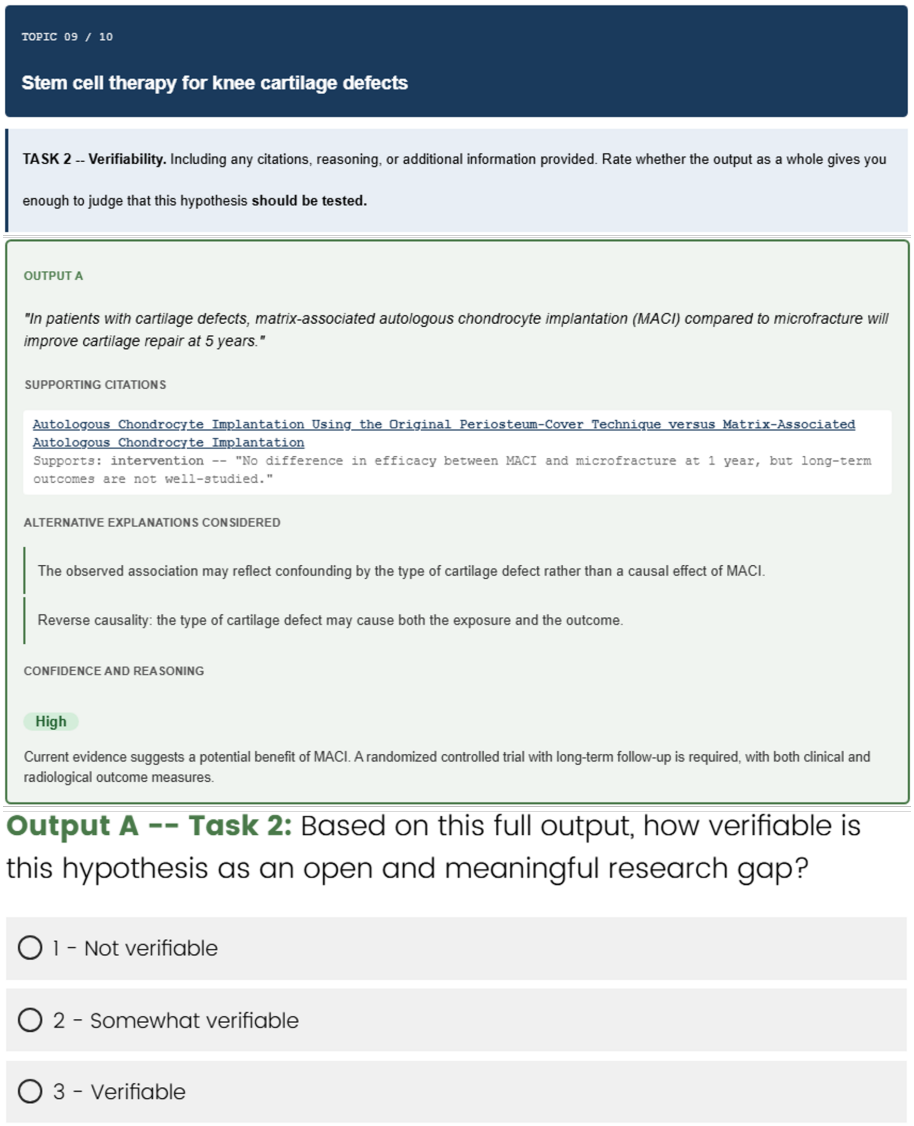

Figure 11: **Hypothesis generation human evaluation. Task 2: Verifiability**, asking whether the hypothesis represents an open and meaningful research gap, given provided citations and competing alternatives. Shown is **EGPC-MEDSCIENTIST output providing structured hypothesis with evidences**.

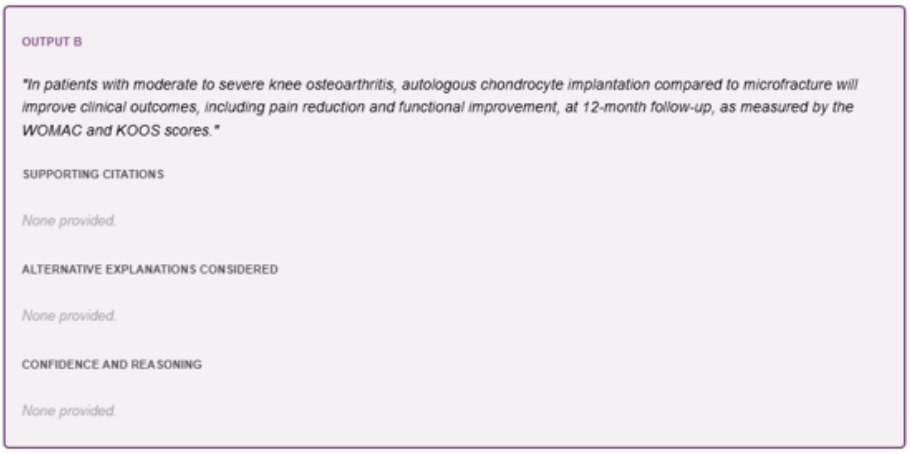

Figure 12: **Hypothesis generation human evaluation. Task 2: Verifiability**, asking whether the hypothesis represents an open and meaningful research gap, given provided citations and competing alternatives. Shown is **Ai-Scientist output doesn't provide verifiable sources**.

First, identify the PICO elements of this paper from its title and abstract

(max 5 words each, write 'not mentioned' if absent):
- Population: who are the study subjects?
- Intervention: what is the main intervention or exposure?
- Comparator: what is the comparison group?
- Outcome: what is the main outcome measured?

Then compare each PICO element against the research conditions above A paper that mismatches on even one PICO element should receive a low score (1-2).

Figure 13: PICO-aware instruction used in Stage 1

You are a biomedical evidence extraction agent. Your task is to extract structured information from a scientific paper following the PICO framework used in evidence-based medicine.

Given the paper title and abstract below, extract the following fields:

- population: who are the study subjects, max 20 words (conditions, age, comorbidities)
- intervention: the main treatment or exposure, max 20 words
- comparison: what it is compared against, max 20 words (or 'none')
- outcome: the primary measured result, max 20 words
- study_design: one of [RCT, cohort study, case-control, cross-sectional, systematic review, meta-analysis, case report]
- grade_level: one of [High, Moderate, Low, Very Low] where High=RCT low bias, Moderate=RCT with limits or cohort, Low=observational, Very Low=case report
- key_findings: one-two sentence max 40 words summarizing the main result
- short_desc: one-two sentence max 40 words describing what this paper is about

Title: {TITLE}
Abstract: {ABSTRACT}

Respond ONLY with a valid JSON object and nothing else.
Example format: {JSON_EXAMPLE}

Figure 14: Stage 2 PICO parser agent prompt. Each paper in the retrieved literature set is parsed into a structured PICO record that is passed to the gap analyzer and hypothesis generator.

You are a biomedical evidence gap analysis agent. You will be given a set of papers already parsed into PICO format. Your task is to identify gaps in the evidence space: combinations of Population, Intervention, Comparison, and Outcome that have not been studied, or have been studied only with weak evidence.

For each gap you identify, classify it using one of these gap types:

- population_extension: the intervention was tested in a narrow population and could be extended
- comparison_gap: no direct comparison between two plausible interventions exists
- outcome_gap: an important outcome has not been measured in existing studies
- dose_gap: only one dose or version of the intervention has been studied
- subgroup_gap: the overall effect is known but response in a specific subgroup is unknown

For each gap, assign a priority (High / Moderate / Low) based on:

- clinical importance of the missing evidence
- feasibility of closing the gap with a study
- how directly the gap follows from the provided papers

Do not include any explanation outside the JSON.

PICO-parsed papers: {PICO_LIST}

Figure 15: Stage 2 gap analyzer agent prompt. The agent reads the PICO-parsed corpus and identifies underexplored evidence configurations, which are passed as prioritized gaps to the hypothesis generator.

You are a biomedical research hypothesis generation agent. Your role is to generate a testable research hypothesis grounded strictly in the provided evidence.

You will receive:

1. A set of input papers in PICO format with their evidence quality
2. A list of identified evidence gaps with priority levels

Step 1 – gap analysis:

- Read all gaps carefully and note their types and priority levels
- Identify which gaps share a common population, mechanism, or intervention – these are candidates to combine
- Identify which gaps are distinct enough to deserve a separate hypothesis
- Randomly pick a gap or combination of gaps and plan 1 hypothesis that addresses it

Step 2 – hypothesis generation rules:

- Generate exactly 1 research hypothesis
- Each hypothesis can address one gap or combine multiple related gaps
- Every claim must be traceable to at least one input paper – do not invent populations, interventions, or outcomes
- Each hypothesis must be specific: include direction of effect (increases, reduces, prevents, delays) and measurable outcome
- The evidence_chain must include at least one entry for each PICO element used in the hypothesis
- supports_pico_element in evidence_chain must be exactly one of the PICO elements defined in hypothesis_pico (population, intervention, comparison, or outcome)

Return your response as a JSON object representing a single hypothesis following this structure: {JSON_FORMAT}. Do not include any text outside the JSON.

Input PICO papers: {PICO_LIST}
Identified gaps: {GAP_LIST}

Figure 16: Stage 2 hypothesis generator prompt. The agent selects a gap or combination of gaps from the gap list and generates one grounded hypothesis with a traceable evidence chain.

You are a biomedical research hypothesis generation agent. Your role is to verify a research hypothesis by identifying competing explanations and assessing confidence, then output a complete revised hypothesis.

You will receive an initial hypothesis with its PICO elements, hypothesis statement, and an evidence chain.

Rules:

- competing_hypotheses must list at least 2 concrete alternative mechanisms or confounders, not generic placeholders
- confidence_reasoning must be 3 sentences: (1) what specific evidence from evidence_chain supports this, (2) what key data is missing, (3) what study design would confirm or refute it
- copy evidence_chain exactly as given – do not add, remove, or modify any entries
- you may revise hypothesis_statement and hypothesis_pico only to make them more specific or precise based on the competing analysis and reasoning
- every claim in the revised hypothesis_statement must be covered by a paper_id already in the evidence_chain – do not introduce claims not traceable to an existing entry

Return your response as a complete JSON object with this structure: {JSON_FORMAT}. Do not include any text outside the JSON.

Input PICO papers: {PICO_LIST}
Initial hypothesis: {HYPOTHESIS_JSON}

Figure 17: Stage 2 hypothesis verifier prompt. The verifier identifies competing explanations and assesses confidence, then returns a revised hypothesis with the evidence chain unchanged.

You are a clinical research reviewer. Review the protocol for three things:

(A) Regulation compliance against the checklists below

(B) Completeness and detail of the protocol stages

(C) Methodological correctness of the protocol stages

For each checklist, report:

- covered_items: item IDs adequately addressed
- missing_items: item IDs applicable but absent or insufficient
- not_applicable_items: item IDs genuinely irrelevant for this study type

Study-type rules:

- SaMD: if no software component, set fda_samd.applicable to false and leave item lists empty.

Determining critical_flags (cover both compliance gaps, completeness/detail gaps, and methodological errors):

- Compliance critical: missing eligibility criteria, missing primary endpoint measurement, missing consent procedure, missing risk identification, missing blinding in an RCT, missing allocation method in an RCT.
- Completeness/detail critical: a stage missing essential steps for its purpose (e.g., intervention stage has no dosing or administration procedure), or an action so vague that a reader must invent information to follow it.
- Methodological critical: a stage that is missing essential steps, has a major logical flaw, or ordering error.
- Non-critical (do not flag): missing secondary endpoints, minor documentation gaps, missing compensation if not applicable.
- Keep critical_flags short – only what genuinely blocks the design.

If critical_flags is non-empty, fill revision_targets: for each flag, name the stage to edit and write a short instruction for what to add or clarify. All revisions must go inside an existing stage's actions or parameters. If critical_flags is empty, set revision_targets to an empty list.

CRITICAL STRUCTURE RULE: critical_flags and revision_targets are TOP-LEVEL fields. Do NOT place them inside spirit, irb, hipaa, or fda_samd sections.

Return a JSON object following the required schema. No text outside the JSON.

SPIRIT checklist ({spirit_n} items): {SPIRIT_ITEMS}
IRB checklist ({irb_n} items): {IRB_ITEMS}
HIPAA checklist ({hipaa_n} items): {HIPAA_ITEMS}
FDA SaMD checklist ({samd_n} items – skip if no software component): {SAMD_ITEMS}

Hypothesis: {HYPOTHESIS_JSON}
Protocol to review: {PROTOCOL_JSON}

Figure 18: Stage 3 compliance reviewer prompt. The reviewer checks regulation compliance, completeness, and methodological correctness, returning critical_flags and revision_targets used by the reviser.

You are a clinical trial protocol designer. Revise the experiment design protocol to address the compliance gaps identified by the regulation reviewer.

Must fix (critical – protocol cannot proceed without these):
{CRITICAL_TARGETS}

Should document (applicable items not yet addressed – add where they fit naturally):
{NON_CRITICAL_MISSING}

Rules:

- For must-fix items, modify only the stages listed above.
- For should-document items, add to the most appropriate existing stage if it fits naturally.
- All additions go into an existing stage's steps or outlines.
- Do NOT add new parts or change the top-level structure.
- Keep all parts that were not targeted – return them unchanged.
- Each new step must be specific and reproducible: include instrument, dose, scale, or timepoint so a reader can follow without guessing.
- Return the COMPLETE protocol JSON with the same top-level structure.

No text outside the JSON.

Research question: {RESEARCH_QUESTION}
Protocol to revise: {PROTOCOL_JSON}

Figure 19: Stage 3 protocol reviser prompt. Given the critical_flags and revision_targets from the reviewer, the reviser updates only the targeted stages while keeping the rest of the protocol unchanged.

You are an impartial scientific judge evaluating a single research hypothesis.

Score it on three criteria (1 = very poor, 5 = excellent):
- novelty: does the hypothesis go beyond directly restating the findings in the provided papers? In biomedical research, incremental advances are expected, but a novel hypothesis proposes a specific comparison, population, or outcome not already answered by the source papers, rather than summarizing what is already known.
- verifiability: are the population, intervention, comparator, and outcome components logically consistent with each other? Can the hypothesis be empirically tested with a realistic study design?
- evidence_grounding: are ALL claims in the hypothesis explicitly grounded in and cited from the provided research background? A claim must include an explicit citation to a specific source paper in the hypothesis text. The judge must not search the background papers to find support for uncited claims. Claims that are uncited or only implicitly supported are considered ungrounded (score 1-2). Claims that include a citation but the cited paper only supports a related claim rather than the exact one are considered partially grounded (score 3).

Research background (source papers): {PAPER_LIST}

Hypothesis: {HYPOTHESIS}

Figure 20: Judger Stage 2 Hypothesis Genaration

You are an impartial clinical research methodology judge evaluating a single study design protocol.

The protocol was generated from the following hypothesis: {HYPOTHESIS_STATEMENT}
Score the protocol on four criteria (1 = very poor, 5 = excellent):
- completeness: are all necessary steps present to achieve the research objective, without missing any primary steps? Do not penalize for missing details within a step (that belongs in detail), only for missing entire steps or sections that are essential for the study to function.
- detail: does each step provide enough specific information to be followed and reproduced without additional assumptions? A step is detailed enough if the reader knows what to do, with what, and when without guessing. Only penalize when a reader would genuinely need to assume or invent information to proceed.
- validity: are the steps factually correct, clinically reasonable, and logically ordered? This includes whether intervention parameters are appropriate for the condition and population, whether design choices are free from factual errors, and whether steps are placed in a reasonable sequence relative to each other.
- compliance: does the protocol address substantive research ethics and regulatory requirements (e.g., IRB, informed consent, data privacy/HIPAA, participant safety monitoring, adverse event reporting), where applicable to the study context? Do NOT penalize for reporting formalities such as trial acronym, formal title format, or trial registration. Use this rubric based on how many substantive domains are entirely absent:
1 = nearly all substantive domains absent
2 = multiple substantive domains entirely absent
3 = one substantive domain entirely absent
4 = no substantive domain absent, only minor individual gaps within covered domains
5 = all applicable requirements addressed with sufficient detail
Protocol: {PROTOCOL_JSON}

Hypothesis: {HYPOTHESIS}

Figure 21: Judger Stage 3 Protocol Design

You are an impartial clinical research methodology judge evaluating the protocol manuscript provided above.

The manuscript was generated from the following research hypothesis: {HYPOTHESIS_STATEMENT}
Score the manuscript on three criteria (1 = very poor, 5 = excellent):
1. evidence_grounding Are the cited papers appropriate for the specific claims they support? Since this is a protocol paper, the hypothesis does not need direct proof. Instead, evaluate whether the cited evidence collectively builds a scientific rationale for why the hypothesis is worth testing. This includes whether citations support the existence and importance of the problem, the biological or mechanistic plausibility of the proposed intervention, and any indirect evidence from related contexts. Do not penalize the manuscript for lacking direct proof of the hypothesis. Do not consider the number of papers cited, their recency, or variety UNLESS they are very significant (e.g., only cite one paper).
2. protocol_completeness Does the manuscript specify all essential elements of a complete clinical trial protocol? Check whether the substantial elements are present and sufficiently described, consistent with standard reporting requirements for interventional trial protocols (e.g. SPIRIT guidelines), such as:
- Research objective and participant eligibility (inclusion/exclusion criteria)
- Randomization and blinding plan (or justification if not applicable)
- Ethical considerations (consent, data privacy, risk-benefit). Do NOT penalize for missing minor administrative details
Penalize heavily (score 1 or 2) for missing or lack of importance elements, such as missing blinding plan in RCT.
3. clarity Is the manuscript readable, internally coherent, and unambiguous? Do NOT penalize for citation format, incomplete bibliographic details (missing journal, volume, page, DOI), or unusual citation key styles. Evaluate prose quality and organization only.

Figure 22: Judger Stage 4 Paper Writing

You are writing the Background section of a clinical trial protocol paper.

You may ONLY cite papers from the evidence lists below. Use format [PAPER_ID] inline. Do NOT introduce any external references or fabricated references. If a claim cannot be supported by the papers below, write it without a citation.

Hypothesis: {HYPOTHESIS_STATEMENT}

PICO:
Population: {PICO_POPULATION}
Intervention: {PICO_INTERVENTION}
Comparison: {PICO_COMPARISON}
Outcome: {PICO_OUTCOME}

Primary papers (used to derive this hypothesis – cite these for specific claims):
{EVIDENCE_PAPERS_PRIMARY}

Secondary papers (from the same systematic review pool – cite for general background):
{EVIDENCE_PAPERS_SECONDARY}

Write the Background section (max 800 words). Summarize what is currently known. Cite each claim with [PAPER_ID].

Figure 23: Stage 4 background section generation prompt.

You are writing the Rationale section of a clinical trial protocol paper.

Hypothesis: {HYPOTHESIS_STATEMENT}

Evidence gaps: {GAPS}

Background section (already written): {BACKGROUND_WRITTEN}

Write the Rationale section (max 300 words). Explain why a new study is needed using the evidence gaps above. State which gap type (population, comparison, outcome, dose, or subgroup) this study addresses and why it is a priority. Do not repeat what the background already says.

Figure 24: Stage 4 rationale section generation prompt.

You are writing the Methods section of a clinical trial protocol paper.

Translate the study protocol below into a well-structured Methods section. Do NOT add, modify, or remove any design decision from the protocol. Do NOT add citations in this section.

Protocol: {PROTOCOL_JSON}

Regulation compliance notes: {COMPLIANCE_SUMMARY}

Write the Methods section covering: study design, participants, intervention, comparator, outcome measures, sample size justification, and statistical analysis plan.

Figure 25: Stage 4 methods section generation prompt.

You are writing the Expected Outcomes section of a clinical trial protocol paper.

Hypothesis: {HYPOTHESIS_STATEMENT}

Protocol stages: {PROTOCOL_STAGES}

Write the Expected Outcomes section (max 300 words). Describe what results are expected and why they matter clinically.

Figure 26: Stage 4 expected outcomes section generation prompt.

You are writing the Limitations section of a clinical trial protocol paper.

Hypothesis: {HYPOTHESIS_STATEMENT}

Protocol: {PROTOCOL_JSON}

Write the Limitations section (max 300 words). Provide honest limitations of this study design without undermining the entire rationale.

Figure 27: Stage 4 limitations section generation prompt.

You are writing the Ethics section of a clinical trial protocol paper.

Regulation flags:
IRB covered: {IRB_COVERED}
IRB missing: {IRB_MISSING}
HIPAA covered: {HIPAA_COVERED}
HIPAA missing: {HIPAA_MISSING}
Critical flags: {CRITICAL_FLAGS}

Helsinki principles to address:

1. Voluntary informed consent
2. Right to withdraw without consequence
3. Risk-benefit justification
4. Data privacy and confidentiality
5. Independent ethics committee approval
6. Priority of participant welfare over scientific interest

Write the Ethics section (max 300 words). Address each Helsinki principle using the regulation flags to fill in study-specific details.

Figure 28: Stage 4 ethics section generation prompt.

You are refining the {SECTION} section of a clinical trial protocol paper.

Full paper draft (for context): {FULL_DRAFT}

Current {SECTION} section: {SECTION_CONTENT}

Improve this section for clarity, coherence, and completeness. Do not change the scientific content or add new information not already present. {CITATION_REMINDER}

Return only the improved section text, with no preamble or explanation.

Figure 29: Stage 4 section refinement prompt. Applied to each section after initial generation, using the full draft as context to ensure coherence across sections.

You are a clinical trial protocol reviewer checking against the SPIRIT checklist.

Draft manuscript: {FULL_MANUSCRIPT}

SPIRIT checklist items: {SPIRIT_ITEMS}

For each item, determine: covered / partial / missing. If covered or partial, cite the specific sentence or phrase in the manuscript that addresses it.

Return a JSON object:
```
{"spirit_coverage": [{"id": "spirit_9a", "status": "covered", "evidence": "..."}, ...],
"critical_missing": ["spirit_14a"], "coverage_score": 0.73}
```

Figure 30: Stage 4 SPIRIT checklist check prompt. Returns coverage status per item and a coverage score used to decide whether a revision pass is needed.

You are revising a clinical trial protocol manuscript to address missing SPIRIT checklist items.

Missing critical items: {MISSING_ITEMS}

Current manuscript: {FULL_MANUSCRIPT}

Revise only the sections where the missing items can be addressed. Return the complete updated manuscript as a JSON object with all six sections.

Return a JSON object:
```
{"background": "...", "rationale": "...", "methods": "...",
"expected_outcomes": "...", "limitations": "...", "ethics": "..."}
```

Figure 31: Stage 4 SPIRIT revision prompt. Triggered when the SPIRIT check finds critical missing items, revising only the relevant sections.

B.2 Model and LLM-Judger Details

Llama-3.3-70B Configuration. The offline screening model is loaded with 4-bit quantization using the BitsAndBytes library, with NF4 quantization type, double quantization enabled, and bfloat16 compute dtype.

Claude Sonnet 4.6 Configuration. The LLM judger uses the Claude Sonnet 4.6 API with a maximum of 1,024 output tokens per call and a structured JSON output schema to ensure consistent response parsing. The total estimated API cost across all experiments in this work is approximately $130 USD.

B.3 Baselines Adaptation

In this section we explain the approaches that we use to adapt baselines for our evaluation. The main adaptation that applied to all baselines is ensuring all the schemes run using offline model Llama 3.3 70B to ensure has the same cutting knowledge.

Stage 1: BioResearcher We adapt the BioResearcher filter prompt (Luo et al., 2025) as a baseline for Stage 1. BioResearcher requires `target`, `conditions`, and `requirements` in the beginning. We set those values using Claude Sonent 4.6 by decomposing the systematic review inclusion and exclusion criteria provided by (Dennstädt et al., 2024).

Stage 1: LLM-SLR Dennstädt et al. (Dennstädt et al., 2024) evaluate LLM-based title and abstract screening using a Likert scale relevance score. Their prompt asks the model to rate how likely a publication is to be included in a systematic literature review on a scale from 1 (very low probability) to 5 (very high probability), given the title, abstract, and a set of relevant inclusion criteria. The model outputs a single integer score. We evaluate their work in Stage 1 with minimal changes to the output format (from integers to JSON).

Stage 2: AI-Scientist We adapt the idea generation module of AI Scientist (Lu et al., 2026) as a baseline for Stage 2, representing single-agent self-reflection without explicit gap analysis. Three changes are made to fit the biomedical hypothesis generation task. First, the template experiment script is replaced by the set of included papers (title and abstract), providing literature context in place of code context in the original code. Second, the `"Experiment"` field is reinterpreted as `"a testable research hypothesis"`. Third, the Semantic Scholar novelty filter is removed. In their original approach, after generation, each idea passes through a novelty filter that queries Semantic Scholar to determine whether the idea overlaps with existing literature.

Stage 2: BioHypoGen. We adapt BioHypoGen (Qi et al., 2024) as a baseline for Stage 2, representing the multi-agent iterative dialogue paradigm. Two changes are made to fit the offline evaluation setting. First, the single hardcoded background example is replaced by including relevant papers. Papers are formatted as a numbered list of title-abstract pairs, matching the numbered structure of the original system's few-shot examples. All agent prompts, turn-taking logic, and the stopping condition are identical to the original.

Stage 3: BioResearcher. We adapt the experimental design module of BioResearcher (Luo et al., 2025) as the baseline for Stage 3. Two changes are made in our adaptation. First, the PDF-extracted structured literature reports that BioResearcher constructs from downloaded full-text papers are replaced by title and truncated abstract from the Stage 1 screened pool, providing the same text-only context used by our proposed system. Second, the Reference Source and GEO public datasets fields are removed from the prompts. The Reference Source field in BioResearcher traces which specific methodological steps within a source protocol document inspired each generated step. Since we provide only title and abstract, which contain no method-level granularity, the field has no valid target and is therefore removed. The GEO datasets field is removed for the same reason that no dataset search is performed in our offline setting. The 3-step prompt chain, the reflection-and-regenerate mechanism, the score threshold, and the sequential Step 3 accumulation are identical to the original.

Stage 4: SafeScientist. We adapt the `Writer` class of SafeScientist (Zhu et al., 2025) as the baseline for Stage 4. Some changes are made to fit the clinical trial protocol writing task. First, the ML paper section set (Abstract, Method, Results, Introduction, Discussion, Related Work) is replaced by six clinical trial protocol sections: background, rationale, methods, expected outcomes, limitations, and ethics. We use similar prompts to capture details for each section in our work and in SafeScientist. Second, the input is replaced by a Stage

2 hypothesis and Stage 3 protocol JSON. There are no experimental results or a completed code run to draw from. The per-section sequential writing structure, full-draft refinement pass with length guard, and the two-round Semantic Scholar citation retrieval and embedding flow are identical to the original. The key distinction from our proposed system is that this baseline retrieves citations dynamically from Semantic Scholar rather than restricting them to the Stage 1 screened pool, and does not include a SPIRIT compliance check or revision loop. Those differences are also reflected in the slight variations in prompts between our work and SafeScientist across sections.

C Additional Results

C.1 PICO Sensitivity in Literature Review (Figure 32).

C.2 Backbone Model Ablation

We re-run EGPC-MEDSCIENTIST replacing Llama-3.3 70B with Qwen-2.5 72B (Yang et al., 2024; Team, 2024)[5] to examine backbone sensitivity (Tables 10). EGPC-MEDSCIENTIST with Qwen still maintains stronger evidence grounding and compliance than Llama-based baselines across stages, indicating that the framework's trustworthiness mechanisms contribute beyond the backbone choice.

		Details ↑	Correctness ↑	Compliance ↑
Stage 2	Llama	**3.00**	**3.20**	**4.40**
	Qwen	2.90	2.80	4.10
Stage 3	Llama	**3.95**	**3.50**	**4.10**
	Qwen	3.20	3.10	3.70
Stage 4	Llama	**3.90**	3.20	**3.30**
	Qwen	3.70	**3.30**	3.00

Table 10: Backbone ablation across stages.

[5] https://hf.co/Qwen/Qwen2.5-72B-Instruct

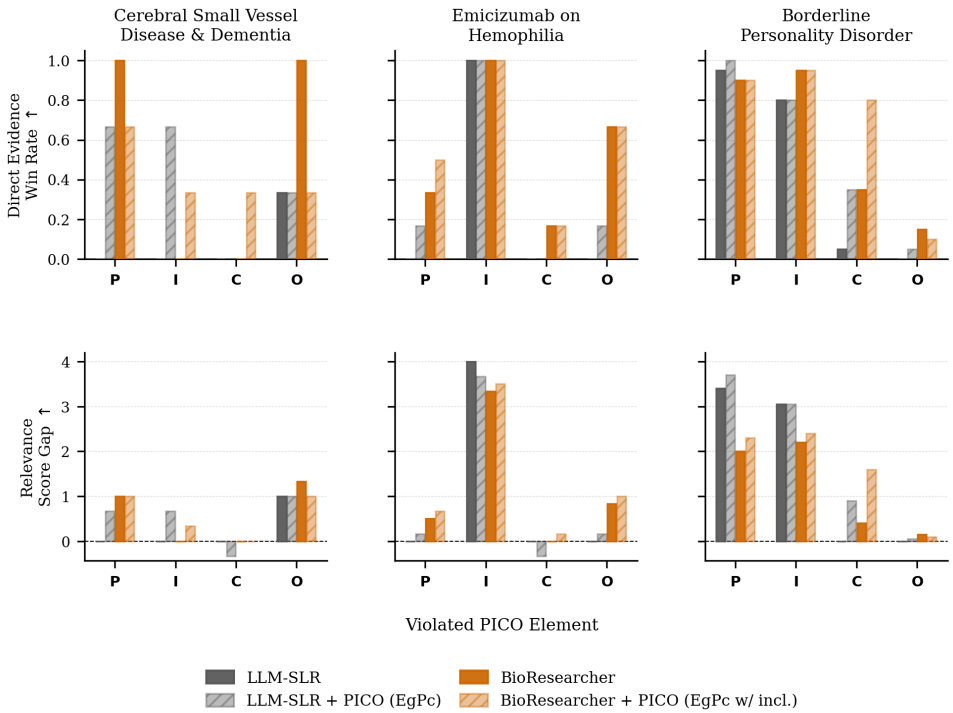

Figure 32: **GRADE-PICO scorer improve baselines** discriminative power over violated PICO element literature even for challenging topic like "Cerebral small vessel disease & Dementia". In this case, all baselines failed to detect any Intervention and Comparison violation.