

# MetaBench: A Multi-task Benchmark for Assessing LLMs in Metabolomics

Yuxing Lu<sup>♠♥</sup>, Xukai Zhao<sup>♠</sup>, J. Ben Tamo<sup>♠</sup>, Micky C. Nnamdi<sup>♠</sup>, Rui Peng<sup>♥</sup>,  
Shuang Zeng<sup>♥</sup>, Xingyu Hu<sup>♠</sup>, Jinzhuo Wang<sup>†♥</sup>, and May D. Wang<sup>†♠</sup>

<sup>♠</sup>Georgia Institute of Technology

<sup>♥</sup>Peking University <sup>♣</sup>Tsinghua University

<sup>†</sup>Corresponding authors: wangjinzhuo@pku.edu.cn, maywang@gatech.edu

## Abstract

Large Language Models (LLMs) have demonstrated remarkable capabilities on general text; however, their proficiency in specialized scientific domains that require deep, interconnected knowledge remains largely uncharacterized. Metabolomics presents unique challenges with its complex biochemical pathways, heterogeneous identifier systems, and fragmented databases. To systematically evaluate LLM capabilities in this domain, we introduce **MetaBench**, the first benchmark for metabolomics assessment. Curated from authoritative public resources, MetaBench evaluates five capabilities essential for metabolomics research: *knowledge*, *understanding*, *grounding*, *reasoning*, and *research*. Our evaluation of 25 open- and closed-source LLMs reveals distinct performance patterns across metabolomics tasks: while models perform well on text generation tasks, cross-database identifier grounding remains challenging even with retrieval augmentation. Model performance also decreases on long-tail metabolites with sparse annotations. With MetaBench, we provide essential infrastructure for developing and evaluating metabolomics AI systems, enabling systematic progress toward reliable computational tools for metabolomics research.

## 1 Introduction

Large Language Models (LLMs) are being rapidly adopted across metabolomics research, driven by their demonstrated success in adjacent biomedical domains (Liu et al., 2025; Bekbergenova et al., 2025). Biomedical LLMs have transformed protein structure prediction, clinically assist with diagnosis and treatment planning, and chemistry-focused systems support reaction prediction and molecular design (Wang et al., 2023; Zhao et al., 2023; Hu and Lu, 2024). Research groups now routinely utilize LLMs for tasks ranging from analyzing experiment results to generating study proposals. However,

this rapid adoption has outpaced systematic evaluation: we lack a comprehensive understanding of which metabolomics tasks LLMs can reliably perform, where they fail, and why. This evaluation gap poses significant risks for a field where incorrect metabolite assignments or pathway interpretations can propagate through analysis pipelines and lead to false biological conclusions.

The consequences of deploying LLMs without proper evaluation extend beyond individual research errors. Metabolomics research demands capabilities that differ fundamentally from general text understanding and generation (Bifarin et al., 2025). Researchers must integrate information across specialized databases such as the Human Metabolome Database (HMDB) and the Kyoto Encyclopedia of Genes and Genomes (KEGG) (Wishart et al., 2022; Kanehisa et al., 2025), each employing distinct identifier systems and ontologies. Without knowing which of the capabilities current LLMs possess, researchers cannot make informed decisions about where to deploy these tools, what verification procedures to implement, or how to interpret AI-assisted results.

Current biomedical evaluation frameworks cannot answer these questions (Chen et al., 2025; Krithara et al., 2023; Jin et al., 2019). These benchmarks evaluate capabilities in natural language understanding but do not measure the specialized operations that metabolomics requires. For example, high performance on MedQA provides no evidence for reliability on identifier grounding, and failure on BioASQ does not preclude success on pathway description generation. Without evaluation frameworks designed for metabolomics-specific tasks, the field lacks criteria for model selection, failure mode identification, or systematic improvement.

In this paper, we present **MetaBench**<sup>1</sup>, a comprehensive benchmark designed to systematically

<sup>1</sup><https://huggingface.co/datasets/LuYuxing/MetaBench>

assess LLM capabilities across metabolomics tasks. MetaBench evaluates five capability levels essential for metabolomics research: *Knowledge* (factual recall of metabolite properties), *Understanding* (generation of coherent pathway descriptions), *Grounding* (accurate identifier mapping across heterogeneous databases), *Reasoning* (extraction of structured relationships from natural language), and *Research* (synthesis of comprehensive study descriptions). Through evaluation of 25 state-of-the-art models on ~8,000 test cases derived from authoritative resources including HMDB, KEGG, PathBank, MetaKG, and MetaboLights (Wishart et al., 2022; Guijas et al., 2018; Lu, 2025; Yurekten et al., 2024), we provide the first systematic characterization of LLM performance in metabolomics, revealing which capabilities current models possess, where critical bottlenecks exist, and what architectural innovations are needed. This work makes the following contributions:

- We introduce MetaBench, the first comprehensive benchmark for evaluating LLM capabilities in metabolomics, comprising ~8,000 test cases across five core capability levels.
- We conduct a systematic evaluation of 25 LLMs, revealing how these models perform across different metabolomics tasks.
- We identify and analyze critical bottlenecks in current LLMs for metabolomics applications, providing mechanistic insights into failure modes and pathways for improvement.

## 2 Related Work

### 2.1 Benchmarks in Scientific Domains

The rapid advancement of LLMs has spurred the development of numerous benchmarks to assess their capabilities in the scientific domain, largely concentrating on broad domains like biomedicine, urban planning, psychology, and chemistry (Cai et al., 2024; Ren et al., 2024; Zhao et al., 2025). Biomedical benchmarks have traditionally focused on tasks like question answering over literature and named entity recognition from clinical notes (Krithara et al., 2023; Jin et al., 2019; Jiang et al., 2025a). While invaluable, these benchmarks do not capture the domain knowledge, data structures, and specific tasks central to metabolomics, such as biochemical pathways, study proposals, and metabolite identifier groundings.

Similarly, in chemistry, benchmarks like MoleculeNet (Wu et al., 2018) have emerged to eval-

uate models on tasks such as molecule property prediction and name conversion. These benchmarks are tailored to the conventions of chemistry, focusing on chemical structure representations (e.g., SMILES) and reaction kinetics, which only partially overlap with the challenges in metabolomics (Roessner and Bowne, 2009). Metabolomics requires an understanding not just of individual molecules but of their roles within complex, dynamic biological systems (Weckwerth, 2003). Our MetaBench is the first to bridge this gap, offering a suite of tasks designed specifically to test the deep, systems-level capabilities required for metabolomics.

### 2.2 NLP in Metabolomics

The application of Natural Language Processing (NLP) in metabolomics is an emerging but critical area of research (Bifarin et al., 2025; Coler et al., 2024). Prior work has primarily focused on leveraging NLP to aid in the interpretation of experimental results by contextualizing findings with existing literature or databases (Lu, 2025; Bekbergenova et al., 2025; Rahman et al., 2024). These efforts often involve tools for named entity recognition (NER) to identify metabolites in text and relation extraction to link them to diseases or genes. However, these applications typically operate as pipeline components rather than end-to-end reasoning systems.

Furthermore, a significant challenge in the metabolomics field is knowledge fragmentation; information about a single metabolite may be spread across multiple databases (e.g., KEGG (Kanehisa et al., 2025), HMDB (Wishart et al., 2022), PubChem (Canese and Weis, 2013)), each using different identifier systems. This necessitates robust entity grounding to create a unified understanding (Ji et al., 2020; Weston et al., 2019). While entity grounding is a well-established NLP task, its application to the diverse and often overlapping identifiers in metabolomics is a unique and unsolved challenge. MetaBench formalizes these challenges into benchmark tasks, positioning our work as a necessary next step: moving from evaluating specific, isolated NLP tasks to a holistic assessment of powerful, end-to-end generative language models.

## 3 MetaBench

### 3.1 Overview

The lack of standardized evaluation benchmarks in metabolomics has left the field without clear crite-

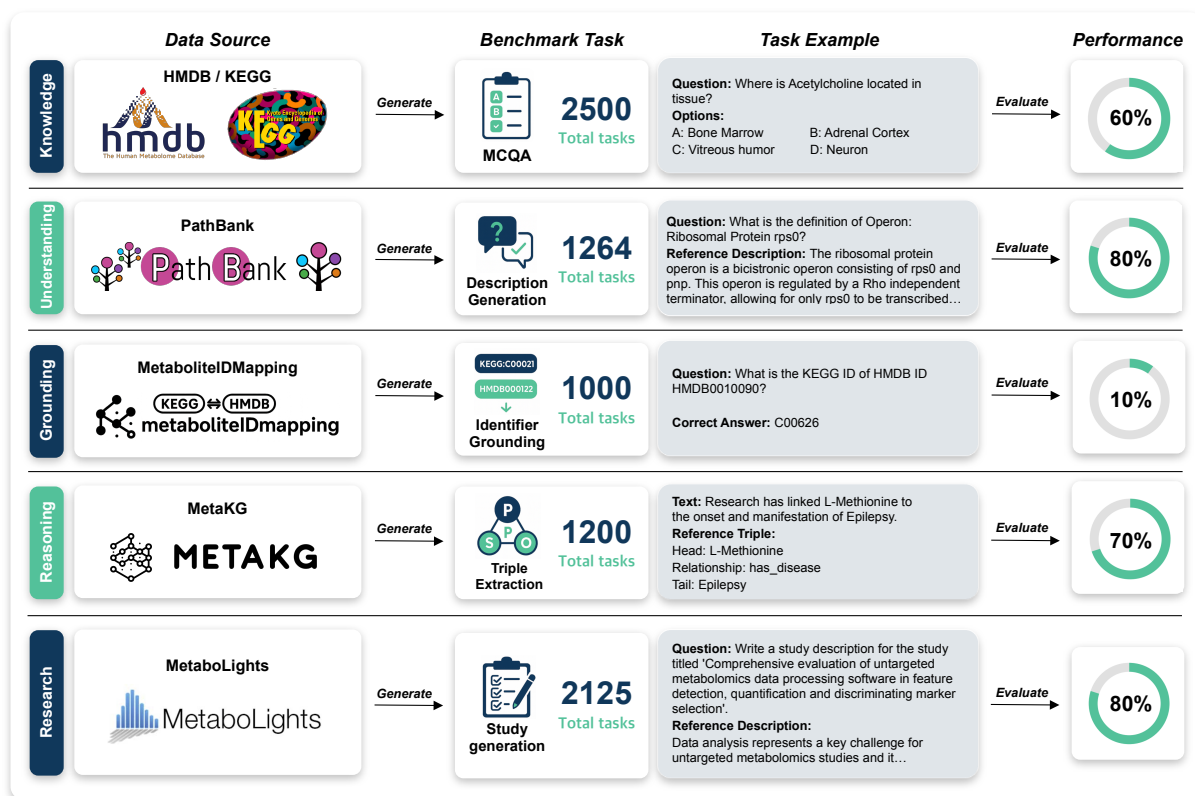


Figure 1: MetaBench construction and task design. MetaBench integrates data from multiple datasets to assess five capabilities: knowledge, understanding, grounding, reasoning, and research.

ria to assess LLM performance, making it difficult to compare methods, identify limitations, or provide reliable guidance for practical use. This gap slows both scientific progress and the safe adoption of LLMs in metabolomics research. To address this challenge, we curated a comprehensive benchmark from different sources (detailed introduction in **Appendix A**) that evaluates LLMs across five capability levels, from factual knowledge recall to research-oriented text generation (**Figure 1**). This benchmark establishes the first critical foundation for methodological innovation and future applications of LLMs in metabolomics.

### 3.2 Benchmark Taxonomy

We organize MetaBench into five capabilities, each aligned with a specialized requirement in metabolomics workflows:

**Knowledge.** Factual recall of general knowledge, such as the correct metabolite taxonomy.

**Understanding.** Generation of coherent descriptions of certain metabolites and pathways;

**Grounding.** Accurate alignment of identifiers across biomedical databases;

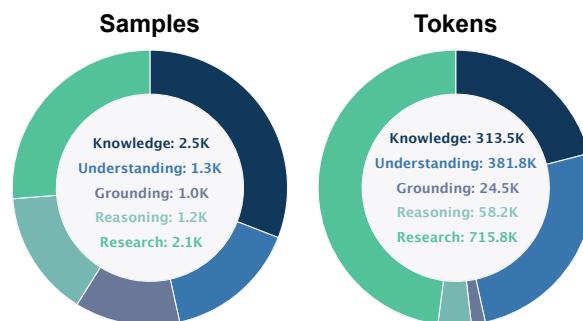


Figure 2: The statistics of MetaBench datasets.

**Reasoning.** Entity and relationship extraction and structuring from natural language;

**Research.** Generation of research-related study descriptions from minimal prompts.

Together, these datasets provide a quantifiable, hierarchical framework for assessing LLMs in metabolomics, analogous to capability ladders used in general NLP benchmarks (Jiang et al., 2025b).

### 3.3 Dataset Construction

We instantiate these capabilities through public available datasets like HMDB (Wishart et al., 2022), KEGG (Kanehisa et al., 2025), MetaboLights (Yurekten et al., 2024), PathBank (Wishart

et al., 2024), and MetaKG (Lu, 2025), the overall process is illustrated in Figure 1. All datasets are released publicly on HuggingFace<sup>2</sup> for reproducibility and the development of LLMs in metabolomics.

### 3.3.1 Knowledge-based MCQA

To assess how well different LLMs encode metabolomics knowledge, we construct a multiple-choice QA benchmark derived from HMDB (Wishart et al., 2022) and KEGG (Kanehisa et al., 2025). We select 26 attributes from these databases spanning taxonomy, molecular properties, biological associations, and pathway relationships. For each attribute, we extract entry-attribute pairs and generate four-option questions by combining the correct value with three in-domain distractors randomly sampled from other values of the same attribute. All questions follow standardized templates (Appendix Table 3), such as "What is the taxonomy class of subject?" The resulting dataset comprises 2,500 questions, with 100 questions per attribute. Examples are provided in Appendix C.

### 3.3.2 Description generation

Beyond factual knowledge assessment, we evaluate whether LLMs can generate coherent, scientifically accurate descriptions of metabolomics concepts. This task tests the model’s ability to produce informative pathway descriptions from pathway names alone. We curated a benchmark using 1,264 pathway-description pairs from PathBank (Wishart et al., 2024), where each pathway name serves as the input prompt and the corresponding expert-written description serves as the reference output. Descriptions range from 148 to 6,452 tokens and contain comprehensive information about biochemical mechanisms, participating enzymes, and biological significance. Performance is assessed using BERTScore to measure semantic similarity between generated and reference texts. Examples are provided in Appendix D.

### 3.3.3 Cross-Database Identifier Grounding

A fundamental challenge in metabolomics research is integrating information about metabolites across heterogeneous databases, which requires accurate mapping between different identifier systems. Only through successful identifier resolution can fragmented knowledge be unified. To evaluate this grounding capability, we

<sup>2</sup>Anonymized for double-blind review

Table 1: Statistics of the MetaBench datasets. Min / Avg. / Max indicates token lengths per sample.

Source	Capability	Min	Avg.	Max
HMDB, KEGG	Knowledge	15	48.17	2113
PathBank	Understanding	26	166.28	880
MetabolitesID	Grounding	7	11.94	35
MetaKG	Reasoning	9	17.95	35
MetaboLights	Research	17	222.57	852

curated a benchmark using the MetabolitesID<sup>3</sup> package that requires models to map metabolite identifiers across KEGG, HMDB, and ChEBI databases. The task encompasses bidirectional mappings: KEGG↔HMDB, KEGG↔ChEBI, HMDB↔ChEBI, and identifier↔name conversions. We sample 1,000 mapping pairs from a comprehensive table containing 130,005 entries, ensuring balanced representation across all mapping types. Each question provides a source identifier and requires the model to predict the corresponding target identifier, evaluated using exact match accuracy. Examples are provided in Appendix E.

### 3.3.4 Knowledge extraction reasoning

To evaluate structured knowledge extraction capabilities, we develop a benchmark that tests whether models can accurately parse metabolomics facts from natural language into knowledge graph triples. We select 1,200 triples from MetaKG (Lu, 2025) spanning six relationship types: has\_disease, has\_disposition, has\_smiles, has\_synonym, has\_class, and has\_tissue\_location. For each triple, we use DeepSeek-V3.1 (Liu et al., 2024) to generate natural language sentences expressing the relationship. For example, the triple (*Inosine*, *has\_tissue\_location*, *Platelet*) is rendered as "Within the human body, the blood’s platelets are a known reservoir for the compound inosine." The generation prompt is provided in Appendix F. Models must extract the complete (head, relationship, tail) triple from each sentence, requiring both natural language understanding and structured reasoning. Performance is evaluated using exact match accuracy, where all three components must be correctly identified. Representative examples are shown in Appendix G.

### 3.3.5 Study description generation

The ultimate goal of applying LLMs to metabolomics is to enable comprehensive

<sup>3</sup><https://github.com/yight/metaboliteIDmapping>

Table 2: MetaBench results for open- and closed-source LLMs across five capabilities. Cells report scores on a 0–100 scale. Best per column is shaded pink and second-best green.

Model	Knowledge	Understanding	Grounding	Reasoning	Research	Average
<b>Open-source models</b>						
Qwen3-1.7b	32.01	81.46	0.27	56.08	79.46	49.86
Qwen3-4b	35.94	83.21	0.27	66.39	80.52	53.27
Qwen3-8b	43.58	82.87	0.25	66.47	80.83	54.40
Qwen3-14b	52.62	82.95	0.29	65.14	80.88	56.38
Qwen3-32b	56.42	83.01	0.33	67.72	82.57	58.01
Gemma-3-270m-it	13.61	81.28	0.20	11.17	80.37	37.33
Gemma-3-1b-it	28.89	81.43	0.47	44.04	81.70	47.31
Gemma-3-4b-it	46.26	81.97	0.53	68.50	82.29	55.91
Gemma-3-12b	51.25	82.64	0.60	70.22	82.73	57.49
Gemma-3-27b	55.82	83.05	0.70	72.69	83.24	59.10
Llama-3.2-1b	12.44	82.40	0.37	33.04	81.77	42.80
Llama-3.2-3b	42.93	82.75	0.60	66.83	82.33	55.49
Llama-3.1-8b	50.58	82.71	0.57	71.92	83.25	57.81
Llama-3.1-70b	57.52	83.08	0.63	72.98	83.18	59.88
DeepSeek-v3.1	54.34	82.38	0.48	73.81	83.52	58.91
GPT-oss-20b	50.66	82.83	0.50	70.16	82.87	57.40
<b>Closed-source models</b>						
GPT-4o-mini	53.82	82.36	0.27	66.39	80.02	56.17
GPT-5-nano	35.65	82.71	0.27	68.47	82.00	53.82
GPT-5-mini	59.02	82.95	0.40	71.89	82.39	59.33
GPT-5	60.50	83.52	0.47	72.39	82.72	59.92
Claude-haiku-3.5	44.34	82.15	0.47	70.64	81.41	55.80
Claude-sonnet-3.7	59.78	82.87	0.53	71.31	83.87	59.67
Claude-sonnet-4	60.94	83.20	0.87	72.56	83.39	60.99
Gemini-2.5-flash	52.10	82.37	0.37	69.25	82.43	57.30
Gemini-2.5-pro	59.54	83.12	0.63	73.14	83.28	60.34

Higher is better for all columns.

research support, from interpreting experimental results through multi-agent systems and data analysis pipelines to potentially designing or executing experiments autonomously. As a step toward this vision, we evaluate models on a research-level task requiring the generation of detailed study descriptions from concise titles. We curated a benchmark using 2,125 title-description pairs from MetaboLights (Yurekten et al., 2024), a leading metabolomics data repository. Each study title serves as the input prompt, and models must generate the corresponding comprehensive description containing experimental methodologies, analytical techniques, key findings, and biological implications. Examples are provided in [Appendix H](#).

### 3.4 Statistics

**Table 1** and **Figure 2** present comprehensive statistics for the MetaBench datasets. We report minimum, average, and maximum token lengths per sample across all tasks. MetaBench comprises

8,100 samples distributed across five capability levels, with Knowledge (2,500 samples) and Research (2,125 samples) representing the largest subsets. Token volume varies substantially by task type: Understanding and Research tasks generate the highest token counts due to paragraph-level outputs (averaging 166 and 223 tokens, respectively), while Grounding and Reasoning tasks involve shorter structured responses (averaging 12 and 18 tokens). This distribution ensures comprehensive evaluation across diverse task formats, from concise factual retrieval and structured extraction to extended scientific text generation.

### 3.5 Evaluation

We employ task-appropriate metrics to ensure rigorous and meaningful evaluation. For classification tasks (knowledge MCQA, identifier grounding, and triple extraction reasoning), we report exact match accuracy. For generation tasks (pathway description generation and study description), we use BERTScore (Zhang et al., 2019) (RoBERTa (Liu

et al., 2019) as backbone model) to measure semantic similarity between generated and reference texts. All systems prompts for each task are provided in **Appendix J**.

For closed-source models, we perform inference through official APIs provided by OpenAI, Anthropic, and Google. For open-source models, we deploy them locally on H200 GPU clusters using the vLLM<sup>4</sup> framework for efficient inference. We standardize all inference settings: temperature is set to 0.1 to encourage deterministic outputs, maximum generation length is capped at 4,096 tokens, and thinking modes are disabled for fair comparison. Each task uses a tailored system prompt that specifies the expected output format and evaluation criteria, ensuring alignment between model responses and metric requirements.

## 4 Results

### 4.1 Overall performance across capabilities

We evaluate 25 state-of-the-art LLMs spanning eight model families (**Appendix Table 4**). **Open-source models:** *Qwen3* (Yang et al., 2025) (1.7B, 4B, 8B, 14B, 32B), *Gemma-3* (Team et al., 2025) (270M-it, 1B-it, 4B-it, 12B, 27B), *Llama* (Dubey et al., 2024) (3.2-1B, 3.2-3B, 3.1-8B, 3.1-70B), *DeepSeek-v3.1* (Liu et al., 2024), and *GPT-oss-20b* (Agarwal et al., 2025). **Closed-source models:** *GPT* (Achiam et al., 2023; Hurst et al., 2024) (4o-mini, 5-nano, 5-mini, 5), *Claude* (Anthropic, 2024, 2025) (haiku-3.5, sonnet-3.7, sonnet-4), and *Gemini-2.5* (Naveed et al., 2025) (flash, pro). All evaluations use standard inference settings without external tools unless explicitly noted.

**Table 2** presents performance across five capabilities. Overall, *Claude-sonnet-4* achieves the highest average score (60.99), followed closely by *Gemini-2.5-pro* (60.34) and *GPT-5* (59.92). Among open-source models, *Llama-3.1-70b* leads at 59.88, nearly matching the best closed-source systems, followed by *Gemma-3-27b* (59.10), *DeepSeek-v3.1* (58.91), *Qwen3-32b* (58.01), and *GPT-oss-20b* (57.40). The narrow performance gap between top open-source and closed-source models (less than 1.5 points) demonstrates that open models can achieve competitive metabolomics reasoning when properly scaled.

Task-specific analysis reveals distinct capability profiles. For Knowledge tasks, *Claude-sonnet-4* marginally leads at 60.94, slightly ahead of *GPT-5*

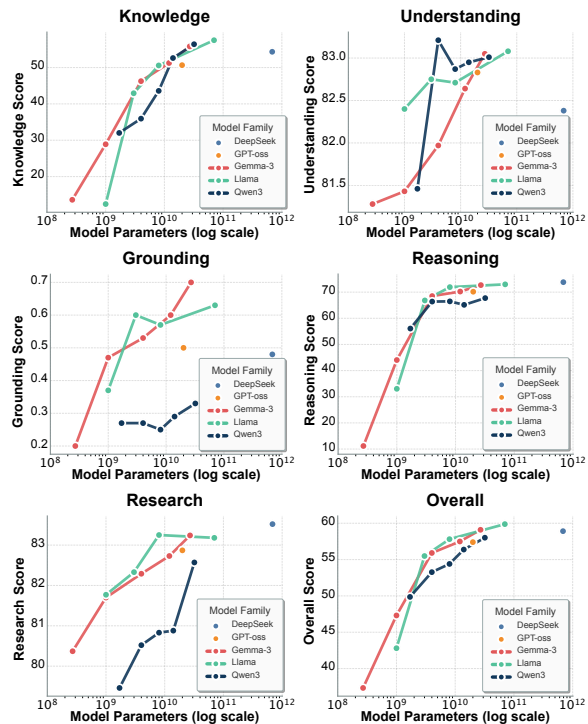


Figure 3: Parameter scaling trends in MetaBench. Average performance increases roughly log-linearly with model size, with consistent family-wise improvements.

(60.50) and *Gemini-2.5-pro* (59.54). Understanding shows remarkably compressed performance, with all competitive models clustering between 81–84%; *GPT-5* achieves the highest score (83.52), but the narrow range suggests this capability saturates early across modern architectures. Grounding presents a stark contrast: without retrieval augmentation, even the best model (*Claude-sonnet-4*) achieves only 0.87% accuracy, two orders of magnitude below other tasks. We analyze this bottleneck in detail in §4.3. For Reasoning, *DeepSeek-v3.1* excels at 73.81%, with *Gemini-2.5-pro* (73.14) and *Claude-sonnet-4* (72.56) close behind. Finally, Research returns to a high, compressed band led by *Claude-sonnet-3.7* (83.87), *DeepSeek-v3.1* (83.52), and *Claude-sonnet-4* (83.39).

### 4.2 Parameter scaling trends

The scaling law works on the five metabolomic capabilities, but the slope is task dependent (**Figure 3**, **Table 2**). Averaged across tasks, scores rise roughly with diminishing returns at the largest scales; family leaders at similar sizes differ by only 1–3 points, highlighting the influence of pretraining data and objectives beyond size alone. Family-wise scaling trends are consistent: *Qwen3* improves from 49.86 (1.7B) to 58.01 (32B), *Gemma-3* from

<sup>4</sup><https://docs.vllm.ai/en/latest/>

② What is the KEGG ID for the compound with HMDB ID HMDB0004148?

C13691 ✓

	Model	Output
Direct Generation	Gemini-2.5-flash	C00092
	Gemini-2.5-pro	C02500
	GPT-4o-mini	C00245
	GPT-5-nano	C00031
	GPT-5-mini	C01183
	GPT-5	C05544
	DeepSeek-v3.1	C00026
	Claude-3.5-haiku	C00186
	Claude-3.7-sonnet	C00092
	Claude-4-sonnet	C05984
Generation with Search API	GPT-4o-mini	C13691
	GPT-5	C13691
	Claude-haiku-3.5	C13691
	Claude-sonnet-3.7	C13691
	Claude-sonnet-4	C13691

Figure 4: LLMs consistently fail in metabolite identifier grounding. 11 LLMs were asked to retrieve KEGG ID for HMDB0004148. None produced the correct answer. In contrast, when augmented with web search API, all successfully provided the correct mapping.

37.33 (270M-it) to 59.10 (27B), and *Llama* from 42.80 (3.2-1B) to 59.88 (3.1-70B). *GPT-20b-oss* and *DeepSeek-v3.1* also demonstrate performance consistent with their parameter counts.

Task-level analysis shows distinct behaviors. Knowledge scales cleanly with size, with *Llama* rising from 12.44 to 57.52 and *Qwen3* from 32.01 to 56.42. Reasoning also benefits strongly, where mid-size models plateau in the high 60s and the largest variants approach 73–74. By contrast, Understanding and Research tasks saturate early: most models fall in a narrow 81–84 band, with larger variants adding less than 2 points, suggesting these generation tasks depend more on instruction tuning than sheer scale. Finally, Grounding remains flat; its accuracy does not exceed 0.87% even for the largest models, showing that parameter growth alone cannot resolve identifier mapping. These results underline that while scale reliably improves recall and reasoning, it delivers marginal gains on fluent generation already near ceiling and fails on grounding tasks unless combined with retrieval and schema-aware normalization.

### 4.3 Identifier grounding is the bottleneck

Among all tested capabilities, Grounding is the lowest-performing capability. Without retrieval, accuracy remains near zero and does not exceed 0.87% even for the largest models (Table 2). A probe illustrates the failure: asked “What is the KEGG ID for HMDB0004148?”, 11 LLMs answered incorrectly (e.g., C00092, C02500); none

returned the correct C13691 (Figure 4). With web search API, all five tested systems produced C13691 on this item (Figure 4).

Aggregate results show the same pattern. Across five representative models, accuracy rises from {0.27, 0.47, 0.53, 0.87}% without search to {20.30, 27.77, 32.93, 38.00, 40.93}% with search (Figure 5). This 40–150× gain remains far from ceiling, indicating missing external evidence and normalization rather than decoding issues.




Figure 5: Without search capabilities (blue bars), all models achieve less than 1% accuracy on cross-database metabolite mapping. Enabling web search (orange bars) improves accuracy by 40–150×, though performance remains below 41% even for the best model (GPT-5), indicating that retrieval alone is insufficient to resolve the grounding bottleneck.

Grounding failures may be because the task combines sparse signals, lossy tokenization, and a misaligned objective under moving targets and ambiguous nomenclature. Metabolite IDs are rare in pretraining corpora and often confined to supplementary tables, so models learn weak parametric associations. Subword tokenizers fracture strings such as HMDB0004148 into pieces (e.g., HMD, B000, 4148), undermining exact matching. Next-token training optimizes plausibility rather than exact resolution, yielding confident but incorrect IDs. Concurrent database updates in HMDB and KEGG further skew memorized mappings. Added to this, synonyms, tautomers, salts, stereochemistry, and organism or compartment context create many-to-many name–ID relations that text-only models cannot reliably disambiguate. Retrieval with schema-aware normalization targets these causes; scaling alone does not.

### 4.4 Long tail problem in metabolomics

Metabolite databases such as HMDB exhibit information concentration: metabolites with lower




Figure 6: Long-tail distribution in metabolite knowledge. Attribute density (green line) and model accuracy (blue bars) both decline across HMDB ID ranges.

HMDB IDs (typically discovered earlier) contain substantially more annotated attributes, while those with higher IDs show progressively sparser information (Xia and Sun, 2022). This pattern reflects the field’s trajectory from well-studied central metabolites toward incompletely characterized peripheral compounds.

To quantify how this heterogeneity affects model performance, we stratified metabolites by HMDB ID ranges and measured average attribute density: <100 (236.73 attributes/metabolite), 100-1,000 (161.92), 1,000-10,000 (128.53), 10,000-100,000 (96.30), and >100,000 (63.38). We sampled 200 MCQA questions per bin and evaluated multiple LLMs. Results from *GPT-oss-20b* are shown in **Figure 6**, where accuracy declines monotonically with decreasing density: 63.5% (127/200) for IDs <100, 58.0% (116/200) for IDs between 100 and 1000, 50.5% for IDs between 1000 and 10000, 47.7% for IDs between 10000 and 100000, versus 40.0% (72/180) for IDs >100,000, resulting in a 23.5 percentage gap.

This long-tail effect reveals a fundamental challenge: models perform acceptably on well-studied metabolites dominating training corpora but fail disproportionately on the sparsely annotated compounds that constitute the metabolome’s majority. Simply scaling training data cannot address the problem since the long tail reflects actual knowledge gaps in the field. Addressing this limitation may require active learning strategies prioritizing difficult cases, multi-modal models leveraging structural information when annotations are sparse, or uncertainty-aware systems that flag low-confidence predictions for long-tail metabolites (Wang et al., 2024; Kandpal et al., 2023).

## 5 Discussion

Our evaluation of 25 LLMs reveals a striking capability heterogeneity that challenges assumptions

about LLM readiness for metabolomics. These results indicate that current training paradigms, optimized for fluent text generation, produce models that excel at synthesizing coherent scientific narratives while struggling with precise factual retrieval and structured knowledge operations (Kalai et al., 2025). Most critically, the catastrophic failure on identifier Grounding (<1% without retrieval, 41% maximum with search API) represents a fundamental limitation rather than a data scarcity problem (more emphasis on the grounding tasks can refer to **Appendix K**). The grounding bottleneck has immediate practical implications: metabolomics applications requiring cross-database integration cannot rely on LLM capabilities alone and must implement specialized identifier resolution systems with schema-aware normalization and chemical structure reasoning.

Beyond diagnosing current limitations, MetaBench establishes a framework for targeted model improvement and responsible deployment in metabolomics. Importantly, our benchmark dataset construction method provides a replicable pathway for creating domain-specific fine-tuning (DFT) datasets beyond evaluation. As new models emerge, whether general LLMs with improved scientific reasoning or domain-specific models pretrained on metabolomics corpora, MetaBench enables continuous, standardized assessment beyond aggregate scores that can obscure critical limitations. By providing both evaluation infrastructure and a concrete methodology for dataset construction, MetaBench supports the development of more capable and reliable AI systems for metabolomics research.

## 6 Conclusion

We present MetaBench, the first comprehensive benchmark for evaluating LLMs on metabolomics. Through the evaluation of 25 LLMs across five capabilities, we reveal substantial performance variations across models and demonstrate that while model scaling improves reasoning, current models catastrophically fail on cross-database mapping and long-tail generalization, establishing that metabolomics requires precision and structured knowledge integration beyond current architectures and training corpus. By publicly releasing MetaBench, we provide essential infrastructure for developing scientifically-grounded models capable of supporting real-world metabolomics research

and enabling systematic progress toward reliable AI-assisted discovery.

## Acknowledgments

This research was supported in part through research cyberinfrastructure resources and services, including the AI Makerspace of the College of Engineering (RRID:SCR\_028058), provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA (RRID:SCR\_027619). We also gratefully acknowledge funding and fellowships that contributed to this work, including a Wallace H. Coulter Distinguished Faculty Fellowship, a Petit Institute Faculty Fellowship, and research funding from Amazon and Microsoft Research awarded to Professor May D. Wang.

## 7 Limitation

Since there is currently no publicly available metabolomics-specialized LLM; consequently, MetaBench evaluates general-purpose and broad biomedical models, quantifying cross-domain generalization and establishing a strong baseline for future in-domain systems. The benchmark centers on widely used resources (KEGG, HMDB, ChEBI, PathBank, MetaboLights), choices that favor clean, reproducible comparisons while deferring spectra/structure modalities and additional databases to later releases. We score classification with accuracy and generation with BERTScore to enable scale and consistency. Results are produced under a standardized, tool-free decoding setup, with retrieval analyzed separately, to isolate intrinsic model behavior. These constraints are deliberate and highlight the release’s strengths: clarity, reproducibility, and extensibility, while charting a direct path to add specialized models in future versions.

## 8 Potential risks

While MetaBench establishes evaluation infrastructure for metabolomics AI, we acknowledge considerations for responsible use. Our transparent reporting of performance disparities, such as the identifier grounding challenge and 23.5% variation across metabolite coverage, helps prevent premature deployment while guiding targeted improvements. The modular design enables continuous benchmark evolution as models advance, and our curation from established resources ensures tasks reflect real metabolomics workflows. By providing

granular evaluation across five capabilities rather than single scores, MetaBench enables informed decisions about where LLMs can augment research and where human expertise remains essential, establishing not just current baselines but the measurement framework necessary for systematic advancement of AI in metabolomics.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, and 1 others. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Anthropic. 2024. Model card for claude 3. [https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model\\_Card\\_Claude\\_3.pdf](https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf).
- Anthropic. 2025. System card: Claude opus 4 & claude sonnet 4. <https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf>.
- Madina Bekbergenova, Lucas Pradi, Benjamin Navet, Youssef Taghzouti, Emma Tysinger, Jean-Luc Wolfender, Wout Bittremieux, Fabien Gandon, and Louis-Félix Nothias. 2025. Metabot: Ai-based agent for natural language-based interaction with metabolomics knowledge graphs. In *ISMB/ECCB 2025*.
- Olatomiwa O Bifarin, Varun S Yelluru, Aditya Simhadri, and Facundo M Fernández. 2025. A large language model-powered map of metabolomics research. *Analytical Chemistry*.
- Bioconductor. 2020. metaboliteidmapping: Annotation package with comprehensive mapping of metabolite ids. <https://bioconductor.org/packages/metaboliteIDmapping/>.
- John Braisted and 1 others. 2023. metlinkr: Facilitating metaanalysis of human metabolomics data through automated linking of metabolite identifiers. *Bioinformatics*, 39(12):btad744.
- Hengxing Cai, Xiaochen Cai, Junhan Chang, Sihang Li, Lin Yao, Changxin Wang, Zhifeng Gao, Hongshuai Wang, Yongge Li, Mujie Lin, and 1 others. 2024. Sciassess: Benchmarking llm proficiency in scientific literature analysis. *arXiv preprint arXiv:2403.01976*.
- Tomas Cajka and Oliver Fiehn. 2024. A reproducibility crisis for clinical metabolomics studies. *Trends in Analytical Chemistry*, 168:117332.

- Kathi Canese and Sarah Weis. 2013. Pubmed: the bibliographic database. *The NCBI handbook*, 2(1):2013.
- Qingyu Chen, Yan Hu, Xueqing Peng, Qianqian Xie, Qiao Jin, Aidan Gilson, Maxwell B Singer, Xuguang Ai, Po-Ting Lai, Zhizheng Wang, and 1 others. 2025. Benchmarking large language models for biomedical natural language processing applications and recommendations. *Nature communications*, 16(1):3280.
- Elizabeth A Coler, Wanxuan Chen, Alexey V Melnik, James T Morton, and Alexander A Aksenov. 2024. Metabolomics in the era of artificial intelligence. *Microbiota and Host*, 2(1).
- 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.
- Dmitry Grapov and 1 others. 2018. Consistency, inconsistency, and ambiguity of metabolite names in biochemical databases used for genome-scale metabolic modelling. *Molecular BioSystems*, 14(3):660–667.
- Carlos Guijas, J Rafael Montenegro-Burke, Benedikt Warth, Mary E Spilker, and Gary Siuzdak. 2018. Metabolomics activity screening for identifying metabolites that modulate phenotype. *Nature biotechnology*, 36(4):316–320.
- Yucheng Hu and Yuxing Lu. 2024. Rag and rau: A survey on retrieval-augmented language model in natural language processing. *arXiv preprint arXiv:2404.19543*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Zongcheng Ji, Qiang Wei, and Hua Xu. 2020. Bert-based ranking for biomedical entity normalization. *AMIA Summits on Translational Science Proceedings*, 2020:269.
- Jiyue Jiang, Pengan Chen, Jiuming Wang, Dongchen He, Ziqin Wei, Liang Hong, Licheng Zong, Sheng Wang, Qinze Yu, Zixian Ma, and 1 others. 2025a. Benchmarking large language models on multiple tasks in bioinformatics nlp with prompting. *arXiv preprint arXiv:2503.04013*.
- Zhuohang Jiang, Pangjing Wu, Ziran Liang, Peter Q Chen, Xu Yuan, Ye Jia, Jiancheng Tu, Chen Li, Peter HF Ng, and Qing Li. 2025b. Hibench: Benchmarking llms capability on hierarchical structure reasoning. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 5505–5515.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*.
- Adam Tauman Kalai, Ofir Nachum, Santosh S Vempala, and Edwin Zhang. 2025. Why language models hallucinate. *arXiv preprint arXiv:2509.04664*.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *International conference on machine learning*, pages 15696–15707. PMLR.
- Minoru Kanehisa, Miho Furumichi, Yoko Sato, Yuriko Matsuura, and Mari Ishiguro-Watanabe. 2025. Kegg: biological systems database as a model of the real world. *Nucleic acids research*, 53(D1):D672–D677.
- Christopher A Kretzler, Maria Sorokina, Felina Kretschmer, Daniel H Pieper, and Christoph Steinbeck. 2024. Navigating common pitfalls in metabolite identification and metabolomics bioinformatics. *Metabolomics*, 20(5):104.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. Biosqqa: A manually curated corpus for biomedical question answering. *Scientific Data*, 10(1):170.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Yijiang Liu, Feifan Zhang, Yifei Ge, Qiao Liu, Siyu He, and Xiaotao Shen. 2025. Application of llms/transformer-based models for metabolite annotation in metabolomics. *Health and Metabolism*, pages 7–7.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yuxing Lu. 2025. Knowledge graph and large language model for metabolomics. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 29281–29282.
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2025. A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology*, 16(5):1–72.
- Rex Powers and 1 others. 2022. A checklist for reproducible computational analysis in clinical metabolomics research. *Metabolites*, 12(1):87.
- Maxx Richard Rahman, Ruoxuan Liu, and Wolfgang Maass. 2024. Incorporating metabolic information into llms for anomaly detection in clinical time-series. *arXiv preprint arXiv:2410.12830*.

- Yuanyi Ren, Haoran Ye, Hanjun Fang, Xin Zhang, and Guojie Song. 2024. Valuebench: Towards comprehensively evaluating value orientations and understanding of large language models. *arXiv preprint arXiv:2406.04214*.
- Ute Roessner and Jairus Bowne. 2009. What is metabolomics all about? *Biotechniques*, 46(5):363–365.
- Neil Swainston, Kieran Smallbone, Pedro Mendes, Douglas Kell, and Norman Paton. 2014. Comparative evaluation of open source software for mapping between metabolite identifiers in metabolic network reconstructions: application to recon 2. *Journal of Cheminformatics*, 6(1):1–11.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivièrè, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Benyou Wang, Qianqian Xie, Jiahuan Pei, Zhihong Chen, Prayag Tiwari, Zhao Li, and Jie Fu. 2023. Pre-trained language models in biomedical domain: A systematic survey. *ACM Computing Surveys*, 56(3):1–52.
- Pengkun Wang, Zhe Zhao, HaiBin Wen, Fanfu Wang, Binwu Wang, Qingfu Zhang, and Yang Wang. 2024. Llm-autoda: Large language model-driven automatic data augmentation for long-tailed problems. *Advances in Neural Information Processing Systems*, 37:64915–64941.
- Wolfram Weckwerth. 2003. Metabolomics in systems biology. *Annual review of plant biology*, 54(1):669–689.
- Leigh Weston, Vahe Tshitoyan, John Dagdelen, Olga Kononova, Amalie Trewartha, Kristin A Persson, Gerbrand Ceder, and Anubhav Jain. 2019. Named entity recognition and normalization applied to large-scale information extraction from the materials science literature. *Journal of chemical information and modeling*, 59(9):3692–3702.
- David S Wishart, AnChi Guo, Eponine Oler, Fei Wang, Afia Anjum, Harrison Peters, Raynard Dizon, Zinat Sayeeda, Siyang Tian, Brian L Lee, and 1 others. 2022. Hmdb 5.0: the human metabolome database for 2022. *Nucleic acids research*, 50(D1):D622–D631.
- David S Wishart, Ray Kruger, Aadhavya Sivakumar, Karxena Harford, Selena Sanford, Rahil Doshi, Nitya Khetarpal, Omolola Fatokun, Daphnee Doucet, Ashley Zubkowski, and 1 others. 2024. Pathbank 2.0—the pathway database for model organism metabolomics. *Nucleic acids research*, 52(D1):D654–D662.
- Gert Wohlgemuth, Pradeep Kumar Haldiya, Egon Willighagen, Tobias Kind, and Oliver Fiehn. 2010. The chemical translation service—a web-based tool to improve standardization of metabolomic reports. *Bioinformatics*, 26(20):2647–2648.
- Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530.
- Yinglin Xia and Jun Sun. 2022. *Statistical data analysis of microbiomes and metabolomics*, volume 13. American Chemical Society.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Ozgur Yurekten, Thomas Payne, Noemi Tejera, Felix Xavier Amaladoss, Callum Martin, Mark Williams, and Claire O’Donovan. 2024. Metabolights: open data repository for metabolomics. *Nucleic acids research*, 52(D1):D640–D646.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 1(2).
- Xukai Zhao, He Huang, Tao Yang, Yuxing Lu, Lu Zhang, Ruoyu Wang, Zhengliang Liu, Tianyang Zhong, and Tianming Liu. 2025. Urban planning in the age of large language models: Assessing openai o1’s performance and capabilities across 556 tasks. *Computers, Environment and Urban Systems*, 121:102332.

## A Sources for MetaBench

MetaBench integrates data from six specialized metabolomics resources, each serving distinct roles in the field. This section provides detailed descriptions of these sources and their contributions to our benchmark.

### A.1 Human Metabolome Database (HMDB)

The Human Metabolome Database (HMDB) (Wishart et al., 2022) is the most comprehensive metabolomics database focused on human metabolism, cataloging over 220,000 metabolite entries including both endogenous metabolites and exogenous compounds. Each entry provides extensive biochemical information including molecular descriptors (SMILES, InChI), physicochemical properties, tissue and biofluid locations, disease associations, and literature references using standardized identifiers (HMDB#####). In MetaBench, HMDB serves as a primary source for the Knowledge MCQA task, providing ground truth data for 26 distinct metabolite attributes, and forms one node of the cross-database identifier mapping in the Grounding capability assessment.

### A.2 Kyoto Encyclopedia of Genes and Genomes (KEGG)

The Kyoto Encyclopedia of Genes and Genomes (KEGG) (Kanehisa et al., 2025) integrates genomic, chemical, and systemic functional information, with KEGG COMPOUND containing over 18,000 metabolite entries organized within the context of biochemical pathways, reactions, and molecular interactions. Each compound receives a KEGG identifier (C#####) and is annotated with pathway memberships, enzyme associations (EC numbers), reaction participation, and orthology relationships.

### A.3 PathBank

PathBank (Wishart et al., 2024) is a comprehensive visual database containing over 110,000 detailed pathway diagrams and expert-curated descriptions covering metabolic, signaling, drug action, and disease pathways across multiple organisms, with particular emphasis on human metabolism. Each entry contains a pathway name, comprehensive textual description ranging from concise summaries to extensive multi-paragraph explanations (148–6,452 tokens), interactive diagrams, and lists of participating metabolites and proteins. PathBank serves

as the exclusive source for the Understanding capability assessment in MetaBench, where 1,264 pathway name-description pairs evaluate whether models can generate scientifically accurate, coherent descriptions of biochemical pathways, testing conceptual understanding beyond simple factual recall.

### A.4 MetabolitesIDMapping

MetaboliteIDmapping is a specialized mapping resource providing comprehensive cross-references between metabolite identifiers across major databases, addressing the fundamental challenge of identifier heterogeneity in metabolomics, where the same metabolite may be referenced differently across resources. The mapping table contains over 130,000 verified equivalence relationships linking identifiers from KEGG, HMDB, ChEBI, PubChem, CAS Registry Numbers, and chemical names (IUPAC and common), enabling bidirectional translation between identifier systems. MetabolitesID forms the foundation of the Grounding capability assessment, where 1,000 sampled identifier pairs construct mapping tasks that directly evaluate a critical bottleneck in practical metabolomics workflows: integrating fragmented knowledge across heterogeneous resources.

### A.5 MetaKG

MetaKG (Lu, 2025) is a large-scale metabolomics knowledge graph structuring information from multiple databases into a unified semantic network with over 2 million entities and 10 million relationships, representing the most comprehensive structured representation of metabolomics knowledge to date. The heterogeneous graph contains nodes representing metabolites, diseases, proteins, genes, tissues, and pathways, connected by typed relationships, with each triple derived from authoritative sources and validated for consistency. MetaKG serves as the source for the Reasoning capability assessment, where 1,200 triples spanning six relationship types are converted into natural language sentences, requiring models to parse unstructured text and extract structured (head, relationship, tail) triples, testing both natural language understanding and structured reasoning essential for knowledge extraction from scientific literature.

### A.6 MetaboLights

MetaboLights (Yurekten et al., 2024) is the largest open-access metabolomics data repository hosted

by the European Bioinformatics Institute (EMBL-EBI), containing over 7,000 studies with associated metadata, analytical methods, and raw data files that serve as a comprehensive archive for metabolomics experimental data. Each study (MTBLS####) includes structured submissions with study titles, detailed descriptions, experimental designs, sample information, analytical protocols (LC-MS, GC-MS, NMR), data processing workflows, and expert-written descriptions providing comprehensive information about research objectives, methodologies, key findings, and biological significance. MetaboLights provides the foundation for the Research capability assessment, where 2,125 study title-description pairs evaluate whether models can generate comprehensive, research-level documentation from minimal prompts, testing capabilities essential for advanced metabolomics research support including experimental design synthesis, analytical technique knowledge, and scientific interpretation.

## B MCQA templates

To construct the Knowledge Graph QA dataset, we selected 26 different relations from MetaKG, the curated metabolomics knowledge graph. For each relation, we designed a specific question template in natural language. These templates convert structured triples (*head*, *relation*, *tail*) into interrogative forms that can be used for multiple-choice question answering.

**Table 3** lists the relation types and their associated templates. For example, the relation `has_class` is expressed as “What is the taxonomy class of {subject}?” and `has_disease` as “Which disease is associated with {subject}?”. This template-based construction ensures consistency across questions while covering diverse relation types in metabolomics.

## C Knowledge MCQA examples

The Knowledge capability is assessed through multiple-choice question answering (MCQA) tasks derived from HMDB and KEGG. Below are representative examples across different attribute types, illustrating the diversity of metabolomics knowledge tested in MetaBench. Questions span metabolite taxonomy, molecular properties, biological associations, pathway relationships, and tissue locations. Each question provides four options with one correct answer, where distractors are sampled

from valid values of the same attribute to ensure domain relevance and task difficulty.

### Example 1: Tissue Location

**Question:** Where is Acetylcholine located in tissue?

**Options:**

- A: Bone Marrow
- B: Adrenal Cortex
- C: Vitreous humor
- D: Neuron

**Answer:** D

### Example 2: Taxonomy Classification

**Question:** What is the taxonomy class of 2-Hydroxyestrone?

**Options:**

- A: Aralkylamines
- B: Estrane steroids
- C: Catechols
- D: Very long-chain fatty acids

**Answer:** B

### Example 3: Database Identifier

**Question:** What is the KEGG ID of Tetrahydrobiopterin?

**Options:**

- A: C00864
- B: C00234
- C: C01747
- D: C00272

**Answer:** D

### Example 4: Enzyme Association

**Question:** Which enzyme is associated with Glycine?

**Options:**

- A: 2.6.1.94

Attribute	Question Template
has_class	What is the taxonomy class of {subject}?
has_name	What is the name of {subject}?
has_disease	Which disease is associated with {subject}?
has_substituent	What is a substituent of {subject}?
has_inchi	What is the InChI of {subject}?
has_disposition	What is the disposition of {subject}?
has_uniprot_id	What is the UniProt ID of {subject}?
has_synonym	What is a synonym for {subject}?
has_process	What process is {subject} involved in?
has_enzyme	Which enzyme is associated with {subject}?
related_to_protein	Which protein is {subject} related to?
has_smiles	What is the SMILES notation of {subject}?
has_cellular_location	Where is {subject} located in the cell?
average_molecular_weight	What is the average molecular weight of {subject}?
chemical_formula	What is the chemical formula of {subject}?
related_to_pathway	Which KEGG pathway is {subject} related to?
has_inchikey	What is the InChIKey of {subject}?
has_kegg_id	What is the KEGG ID of {subject}?
has_biospecimen_location	Where is {subject} found as a biospecimen?
has_pathway	What is a KEGG pathway associated with {subject}?
belongs_to_orthology	To which KEGG orthology does {subject} belong?
has_reaction	What is a KEGG reaction involving {subject}?
has_description	What is the description of {subject}?
belongs_to_network	To which KEGG network does Gene {subject} belong?
has_tissue_location	Where is {subject} located in tissue?
is_a_sub_class_of	{subject} is a subclass of what?

Table 3: Relation types and their corresponding question templates.

- B: 2.1.1.156
- C: 1.1.1.228
- D: 4.2.3.87

**Answer: B**

#### Example 5: Pathway Relationship

**Question:** What is a pathway associated with C15976?

**Options:**

- A: SMP00497
- B: SMP00417
- C: SMP00384
- D: SMP00507

**Answer: C**

## D Description generation examples

The Understanding capability is assessed through pathway description generation tasks derived from PathBank. Models receive pathway names as input and must generate comprehensive, scientifically accurate descriptions explaining the biochemical mechanisms, participating molecules, biological significance, and clinical relevance. Below are representative examples demonstrating the range of pathway types and description complexity in MetaBench, from metabolic pathways and disease mechanisms to drug action pathways and gene regulation systems.

#### Example 1: Signaling Pathway

**Question:** What is the definition of Succinate Signalling?

**Reference Description:** Tricarboxylic acid (TCA) cycle intermediates can function as inflammatory signals. Succinate enhances

glycolysis in several ways. It inhibits prolyl hydroxylase domain (PHD) enzyme function, both directly by product inhibition and indirectly via reactive oxygen species (ROS), driving hypoxia-inducible factor-1 $\alpha$  (HIF-1 $\alpha$ ) accumulation and increased glycolysis. ROS also inhibit mitochondrial function, boosting glycolysis as a result. Elevated HIF-1 $\alpha$  enhances the expression of genes containing HIF response elements (HREs), including the interleukin 1 $\beta$  (IL-1 $\beta$ ) gene. Succinate can signal through succinate receptor 1 (SUCNR1) and act in synergy with Toll-like receptors (TLRs) to boost dendritic cell function. NAD<sup>+</sup> exerts several anti-inflammatory effects by activating sirtuins, a class of NAD<sup>+</sup>-dependent deacetylases. Finally, inactivation of Sirt3 has been demonstrated to enhance NLR family, pyrin domain containing 3 (NLRP3) inflammatory activation.

#### Example 2: Gene Regulation

**Question:** What is the definition of Operon: Ribosomal Protein rps0?

**Reference Description:** The ribosomal protein operon is a bicistronic operon consisting of rps0 and pnp. This operon is regulated by a Rho independent terminator, allowing for only rps0 to be transcribed if the terminator is formed.

#### Example 3: Disease Pathway

**Question:** What is the definition of Glycogenesis, Type IA. Von Gierke Disease?

**Reference Description:** Glycogen storage disease type I (also known as GSDI or von Gierke disease) is an inherited disorder caused by the buildup of a complex sugar called glycogen in the body's cells. The accumulation of glycogen in certain organs and tissues, especially the liver, kidneys, and small intestines, impairs their ability to function normally. Researchers have described two types of GSDI, which differ in their signs and symptoms and genetic cause. These types are known as glycogen storage disease type Ia (GSDIa) and glycogen stor-

age disease type Ib (GSDIb). Two other forms of GSDI have been described, and they were originally named types Ic and Id. However, these types are now known to be variations of GSDIb; for this reason, GSDIb is sometimes called GSD type I non-a. Mutations in two genes, G6PC and SLC37A4, cause GSDI.

#### Example 4: Drug Action Pathway

**Question:** What is the definition of Tetracycline Action Pathway?

**Reference Description:** Tetracycline is a short-acting antibiotic that is semi-synthetically produced from chlortetracycline, a compound derived from *Streptomyces aureofaciens*. Tetracycline enters bacterial cells by passively diffusing through membrane porin channels. Once inside the cell, tetracycline reversibly binds to the 30S subunit just above the binding site for aminoacyl tRNA. At its primary binding site, interactions with the sugar phosphate backbone of residues in helices 31 and 34 via hydrogen bonds with oxygen atoms and hydroxyl groups on the hydrophilic side of the tetracycline help anchor the drug in position. Salt bridge interactions between the backbone of 16S rRNA and tetracycline are mediated by a magnesium ion in the binding site. Tetracycline prevents incoming aminoacyl tRNA from binding to the A site on the ribosome-RNA complex via steric hindrance. This causes inhibition of protein synthesis and hence bacterial cell growth.

#### Example 5: Pharmaceutical Mechanism

**Question:** What is the definition of Sufentanil Action Pathway?

**Reference Description:** Sufentanil is a pharmacologically-active synthetic small molecule derived from fentanyl and belongs to a class of drugs called opioids. Opioids are therapeutically employed to achieve analgesia. Sufentanil's rapid mechanism of action primarily involves its agonistic effects on mu-type opioid receptors which are inhibitory G-coupled protein receptors and lead to the inhibition of adenylate cy-

clase and decrease in cAMP production. It also inhibits nociceptive neurotransmitter release and induces membrane hyperpolarization. Analgesia, anesthesia, and respiratory depression are a consequence of remifentanyl's action.

## E Grounding (Id mapping) examples

The Grounding capability is assessed through cross-database identifier mapping tasks derived from MetabolitesID. This task evaluates whether models can accurately translate metabolite identifiers across heterogeneous database systems (KEGG, HMDB, ChEBI) and between structured identifiers and chemical names. Successful identifier grounding is fundamental to integrating fragmented metabolomics knowledge, yet represents a critical bottleneck for current LLMs. Below are representative examples demonstrating the diverse mapping types required in MetaBench.

### Example 1: HMDB to KEGG Mapping

**Question:** What is the KEGG ID of HMDB ID HMDB0010090?

**Answer:** C00626

### Example 2: KEGG to HMDB Mapping

**Question:** What is the HMDB ID of KEGG ID C07251?

**Answer:** HMDB0014982

### Example 3: Name to KEGG Mapping

**Question:** What is the KEGG ID of metabolite Ponalactone A?

**Answer:** C09174

### Example 4: Name to HMDB Mapping

**Question:** What is the HMDB ID of metabolite Ethylparaben?

**Answer:** HMDB0032573

### Example 5: Name to ChEBI Mapping

**Question:** What is the ChEBI ID of metabolite Croconazole?

**Answer:** 3920

**Note:** These examples illustrate the precision re-

quired for identifier grounding tasks. Models must produce exact matches to succeed, as approximate or similar identifiers are incorrect. This task represents the most challenging capability in MetaBench, with even the best-performing models (*Claude-sonnet-4*) achieving only 0.87% accuracy without retrieval augmentation, highlighting a fundamental limitation in current LLM architectures for structured database mapping tasks.

## F Prompt for reasoning benchmark

To construct the Reasoning (Triple Extraction) benchmark, we generated natural language sentences from structured knowledge graph triples using DeepSeek-V3.1 (Liu et al., 2024). Each triple from MetaKG was transformed into a natural language sentence using the prompt template shown below. The prompt explicitly instructs the model to avoid template-like language and to create varied, natural expressions of biomedical facts, ensuring that the resulting benchmark requires genuine semantic parsing rather than pattern matching.

### Prompt for Reasoning Benchmark

Given the following triple from a biomedical knowledge graph:

**Subject:** head

**Predicate:** rel

**Object:** tail

Write a natural language sentence that expresses this fact, but do so in a way that is different from a simple template. Be creative and vary the phrasing. Do not use the words 'subject', 'predicate', or 'object'.

This generation approach ensures that the Reasoning task evaluates genuine natural language understanding rather than template recognition. The diversity in expression styles reflects the variability found in real scientific literature, making the benchmark more representative of practical knowledge extraction scenarios in metabolomics research.

## G Triple extraction examples

The Reasoning capability is assessed through knowledge graph triple extraction tasks derived from MetaKG. Models receive natural language sentences describing metabolite relationships and must extract structured triples in the format (head, relationship, tail). This task evaluates the model's ability to parse unstructured scientific text and

perform structured reasoning, which is a critical skill for automated knowledge extraction from metabolomics literature. Below are representative examples demonstrating disease associations and metabolic disorder relationships.

#### Example 1: Disease Association

**Text:** Research has linked L-Methionine to the onset and manifestation of Epilepsy.

**Extracted Triple:**

- **Head:** L-Methionine
- **Relationship:** has\_disease
- **Tail:** Epilepsy

#### Example 2: Metabolic Disorder

**Text:** Capric acid is implicated in the development of disorders related to metabolism and nutrition.

**Extracted Triple:**

- **Head:** Capric acid
- **Relationship:** has\_disease
- **Tail:** Metabolism and nutrition disorders

#### Example 3: Biomarker Relationship

**Text:** Citrulline serves as a crucial biomarker for the diagnosis and assessment of intestinal failure.

**Extracted Triple:**

- **Head:** Citrulline
- **Relationship:** has\_disease
- **Tail:** Intestinal failure

#### Example 4: Deficiency Manifestation

**Text:** A deficiency in Vitamin A can manifest as a range of afflictions impacting the nervous system.

**Extracted Triple:**

- **Head:** Vitamin A
- **Relationship:** has\_disease
- **Tail:** Nervous system disorders

#### Example 5: Contributing Factor

**Text:** A deficiency in propionylcarnitine can be a contributing factor in the development of missing teeth.

**Extracted Triple:**

- **Head:** Propionylcarnitine
- **Relationship:** has\_disease
- **Tail:** Missing teeth

## H Study description generation examples

The Research capability is assessed through study description generation tasks derived from MetaboLights. Models receive concise study titles as input and must generate comprehensive descriptions containing experimental methodologies, analytical techniques, key findings, and biological implications. This task represents the highest level of scientific complexity in MetaBench, testing the model's ability to synthesize knowledge about research design, technical protocols, and scientific interpretation. Below are representative examples demonstrating diverse research domains in metabolomics.

#### Example 1: Software Evaluation

**Question:** Write a study description for the study titled 'Comprehensive evaluation of untargeted metabolomics data processing software in feature detection, quantification and discriminating marker selection'.

**Reference Description (excerpt):** Data analysis represents a key challenge for untargeted metabolomics studies and it commonly requires extensive processing of more than thousands of metabolite peaks included in raw high-resolution MS data. Although a number of software packages have been developed to facilitate untargeted data processing, they have not been comprehensively scrutinized in the capability of feature detection, quantification and marker selection using a well-defined benchmark sample set. In this study, we acquired a benchmark dataset from standard mixtures consisting of 1100 compounds with specified concentration ratios including 130 compounds with significant variation of concentrations. Five soft-

ware evaluated here (MS-Dial, MZmine 2, XCMS, MarkerView, and Compound Discoverer) showed similar performance in detection of true features... [continues]

#### Example 2: Clinical Metabolomics

**Question:** Write a study description for the study titled 'Altered metabolome and microbiome features provide clues in understanding irritable bowel syndrome and depression comorbidity'.

**Reference Description (excerpt):** Irritable bowel syndrome (IBS) is one of the functional gastrointestinal disorders characterized by chronic and/or recurrent symptoms of abdominal pain and irregular defecation. Changed gut microbiota has been proposed to mediate IBS; however, contradictory results exist, and IBS-specific microbiota, metabolites, and their interactions remain poorly understood. To address this issue, we performed metabolomic and metagenomic profiling of stool and serum samples based on discovery (n = 330) and validation (n = 101) cohorts... [continues]

#### Example 3: Multi-Omics Integration

**Question:** Write a study description for the study titled 'Mechanisms of hepatic steatosis in chickens: integrated analysis of the host genome, molecular phenomes and gut microbiome'.

**Reference Description (excerpt):** Hepatic steatosis is the initial manifestation of abnormal liver functions and often leads to liver diseases such as non-alcoholic fatty liver disease in humans and fatty liver syndrome in animals. In this study, we conducted a comprehensive analysis of a large chicken population consisting of 705 adult hens by combining host genome resequencing, liver transcriptome, proteome and metabolome analysis, as well as microbial 16S rRNA gene sequencing of each gut segment... [continues]

#### Example 4: Developmental Metabolomics

**Question:** Write a study description for the study titled 'Changes in the Milk Metabolome of the Giant Panda (*Ailuropoda melanoleuca*) with Time after Birth – Three Phases in Early Lactation and Progressive Individual Differences'.

**Reference Description (excerpt):** Ursids (bears) in general, and giant pandas in particular, are highly altricial at birth. The components of bear milks and their changes with time may be uniquely adapted to nourish relatively immature neonates, protect them from pathogens, and support the maturation of neonatal digestive physiology. Serial milk samples collected from three giant pandas in early lactation were subjected to untargeted metabolite profiling and multivariate analysis... [continues]

#### Example 5: Subcellular Imaging

**Question:** Write a study description for the study titled 'Subcellular antibiotic visualization reveals a dynamic drug reservoir in infected macrophages'.

**Reference Description (excerpt):** Tuberculosis, caused by the intracellular pathogen *Mycobacterium tuberculosis*, remains the world's deadliest infectious disease. Sterilizing chemotherapy requires at least 6 months of multidrug therapy. Difficulty visualizing the subcellular localization of antibiotics in infected host cells means that it is unclear whether antibiotics penetrate all mycobacteria-containing compartments in the cell. Here, we combined correlated light, electron, and ion microscopy to image the distribution of bedaquiline in infected human macrophages... [continues]

## I Compared LLMs

We evaluate 25 state-of-the-art large language models spanning eight model families to ensure comprehensive coverage of current LLM capabilities in metabolomics. Our selection includes both open-source models, *Qwen3* (1.7B–32B), *Gemma-3* (270M–27B), *Llama* (1B–70B), *DeepSeek-v3.1* (685B MoE), and *GPT-oss-20b* (20B MoE) and

Table 4: Models evaluated in MetaBench. Parameter counts are from official sources where available; proprietary models do not disclose parameter counts.

Model	Params	Architecture	Release date	Provider
<b>Open-source models</b>				
Qwen3-1.7b	1.7B	dense	2025-04-29	Alibaba Qwen
Qwen3-4b	4B	dense	2025-04-29	Alibaba Qwen
Qwen3-8b	8B	dense	2025-04-29	Alibaba Qwen
Qwen3-14b	14B	dense	2025-04-29	Alibaba Qwen
Qwen3-32b	32B	dense	2025-04-29	Alibaba Qwen
Gemma-3-270m-it	270M	dense	2025-03-12	Google
Gemma-3-1b-it	1B	dense	2025-03-12	Google
Gemma-3-4b-it	4B	dense	2025-03-12	Google
Gemma-3-12b	12B	dense	2025-03-12	Google
Gemma-3-27b	27B	dense	2025-03-12	Google
Llama-3.2-1b	1B	dense	2024-09-25	Meta
Llama-3.2-3b	3B	dense	2024-09-25	Meta
Llama-3.1-8b	8B	dense	2024-07-23	Meta
Llama-3.1-70b	70B	dense	2024-07-23	Meta
DeepSeek-v3.1	685B (MoE) <sup>a</sup>	MoE	2025-08-21	DeepSeek
GPT-oss-20b	20B	MoE	2025-08-21	OpenAI
<b>Closed-source models</b>				
GPT-4o-mini	Not disclosed	multimodal	2024-07-18	OpenAI
GPT-5-nano	Not disclosed	unified/thinking variants	2025-08-07	OpenAI
GPT-5-mini	Not disclosed	unified/thinking variants	2025-08-07	OpenAI
GPT-5	Not disclosed	unified/thinking variants	2025-08-07	OpenAI
Claude-haiku-3.5	Not disclosed	hybrid family	2024-10-22	Anthropic
Claude-sonnet-3.7	Not disclosed	hybrid reasoning	2025-02-24	Anthropic
Claude-sonnet-4	Not disclosed	hybrid reasoning	2025-05-22	Anthropic
Gemini-2.5-flash	Not disclosed	thinking model	2025-06-17	Google
Gemini-2.5-pro	Not disclosed	thinking model	2025-03-25	Google

<sup>a</sup> MoE = mixture-of-experts. Count refers to total parameters across experts, not active per token.

closed-source systems from leading providers: OpenAI’s *GPT* series (4o-mini through GPT-5), Anthropic’s *Claude* family (haiku-3.5, sonnet-3.7, sonnet-4), and Google’s *Gemini-2.5* models (flash and pro). This diversity enables systematic analysis of how model scale, architecture (dense vs. mixture-of-experts), and training paradigms affect performance across the five metabolomics capabilities assessed in MetaBench. **Table 4** summarizes the key specifications of all evaluated models, including parameter counts (where available), architectural types, release dates, and providers. This comprehensive model selection establishes MetaBench as a rigorous testbed for tracking progress in metabolomics-oriented language modeling across both academic and industrial research efforts.

## J System prompts

In our evaluation framework, we designed five distinct system prompts, each tailored to a capability

assessed in MetaBench. These prompts ensure that large language models receive clear, domain-specific instructions aligned with metabolomics tasks.

- **MCQA (Multiple-Choice Question Answering):** Tests factual knowledge of metabolites, pathways, enzymes, and biological systems.
- **Triple Extraction:** Assesses reasoning ability by extracting (*head, relationship, tail*) triples from natural language.
- **Study Description:** Evaluates scientific writing and comprehension by generating detailed study descriptions from research titles.
- **QA (Open-ended Question Answering):** Probes explanatory ability on pathways, mechanisms, and processes.
- **Metabolite Mapping:** Measures precision in identifier mapping across metabolomics databases like KEGG, HMDB and ChEBI.

The following parts contain the exact system prompts used for each evaluation task.

### System Prompt for MCQA

**Task:**

Answer multiple choice questions about metabolite properties, pathways, enzymes, and biological processes.

**Instructions:**

- You will be given a question and four options (A, B, C, D).
- Choose the correct answer based on your knowledge of metabolomics, biochemistry, and biological systems.
- Respond with only the letter (A, B, C, or D).

**Coverage:**

- Metabolite taxonomy and classification
- Enzyme associations and EC numbers
- KEGG pathway and network associations
- Molecular properties (weight, structure)
- Tissue locations and biological functions
- Protein and gene relationships

**Examples:**

Q: What is the taxonomy class of Prostaglandin E1?

Options: {A: Prostaglandins and related compounds, B: Phenylpyrazoles, C: Furanones, D: Azolines}

Answer: A

Q: What is the taxonomy class of PC(18:1(9Z)/18:0)?

Options: {A: Alanine and derivatives, B: Glycerophospholipids, C: Anisoles, D: Isoleucine and derivatives}

Answer: B

Q: Where is Acetylcholine located in tissue?

Options: {A: Bone Marrow, B: Adrenal Cortex, C: Vitreous humor, D: Neuron}

Answer: D

### System Prompt for Triple Extraction

**Task:**

Extract knowledge graph triples from natural language text about metabolites and their biological properties.

**Format:**

Head: [metabolite/entity]

Relationship: [relationship type]

Tail: [disease/condition/property/location]

**Allowed relationships:**

- has\\_disease : metabolite associated with diseases/conditions
- has\\_disposition : metabolite location or cellular disposition
- has\\_smiles : molecular structure in SMILES notation
- has\\_synonym : alternative names
- has\\_class : biochemical classification
- has\\_tissue\_location : tissue or organ presence

**Examples:**

Text: Within the human body, the blood's platelets are a known reservoir for the compound inosine.

Head: Inosine

Relationship: has\\_tissue\_location

Tail: Platelet

Text: Dolichyl phosphate D-mannose is categorized among the broader class of carbohydrates and carbohydrate conjugates.

Head: Dolichyl phosphate D-mannose

Relationship: has\\_class

Tail: Carbohydrates and carbohydrate conjugates

Text: The molecular structure of Isobutyryl-L-carnitine is uniquely defined by the SMILES string CC(C)C(=O)O[C@H](CC(O)=O)...

Head: Isobutyryl-L-carnitine

Relationship: has\\_smiles

Tail: CC(C)C(=O)O[C@H](CC(O)=O)...

### System Prompt for Study Description

**Task:**

Write a comprehensive study description for a given metabolomics research title.

**Instructions:**

Include the following:

1. Study purpose and objectives
2. Methodology and experimental design (LC-MS, GC-MS, NMR)
3. Key findings and results
4. Clinical or scientific significance
5. Sample characteristics and analytical methods

6. Metabolic pathways and biological processes involved

**Coverage:**

- Untargeted and targeted metabolomics
- Sample preparation and analytical techniques
- Data processing and statistical analysis
- Biological interpretation and pathway analysis
- Clinical applications and biomarker discovery

**Example:**

Title: "Dissolved organic carbon compounds in deep-sea hydrothermal vent fluids from the East Pacific Rise at ..."

Description: "... detailed description ..."

### System Prompt for QA

**Task:**

Answer questions about biochemical definitions, metabolic pathways, and biological processes.

**Instructions:**

Provide comprehensive, accurate answers including:

1. Definition of the pathway, process, or concept
2. Key enzymes, reactions, and components
3. Biological significance and function
4. Clinical or pathological relevance (if any)
5. Related pathways or mechanisms

**Coverage:**

- Metabolic pathways (TCA cycle, glycolysis, etc.)
- Biochemical processes and mechanisms
- Disease definitions and mechanisms
- Drug action pathways
- Operon regulation
- Molecular signaling

**Examples:**

Q: What is the definition of Succinate Signalling?

A: "... detailed definition ..."

Q: What is the definition of Tetracycline Action Pathway?

A: "... detailed definition ..."

### System Prompt for Metabolite Mapping

**Task:**

Answer metabolite identifier mapping questions between databases and naming conventions.

**Instructions:**

- Provide only the identifier value as the answer.
- No extra explanation or formatting.

**Supported identifiers:**

- CAS Registry Numbers
- PubChem CIDs
- Chemical names (IUPAC, common)
- HMDB IDs
- ChEBI IDs
- KEGG IDs

**Examples:**

Q: What is the KEGG ID of metabolite Tamibarotene?

A: C12864

Q: What is the HMDB ID of metabolite TG(15:0/i-20:0/a-17:0)[rac]?

A: HMDB0102782

Q: What is the ChEBI ID of Benzilic acid?

A: 39414

Q: What is the KEGG ID of 2-Iodophenol?

A: C01874

## K The Critical Role of Metabolite Identifier Grounding in Metabolomics

Metabolite identifier grounding represents a fundamental bottleneck in metabolomics research (Swainston et al., 2014). A single metabolite may be referenced by multiple identifiers: KEGG uses compound IDs (e.g., C00031 for D-Glucose), HMDB employs alphanumeric codes (e.g., HMDB0000122), ChEBI assigns numerical IDs (e.g., 17234), and scientific literature uses IUPAC names, common names, or synonyms, with PubChem listing nearly 250 synonyms for L-Tryptophan alone (Krettler et al., 2024). This heterogeneity arose because different databases emerged from distinct research communities: KEGG emphasizes pathway context, HMDB focuses on human metabolites with clinical relevance, ChEBI provides chemical ontol-

ogy, and PubChem offers comprehensive chemical data (Wishart et al., 2022; Kanehisa et al., 2025). Each adopted identifier schemes suited to their organization, predating standardization efforts.

Accurate identifier grounding is essential throughout the research pipeline. Cross-study meta-analyses depend on harmonizing identifiers across laboratories and platforms (Braisted et al., 2023). Database construction efforts like MetaKG require linking millions of facts from heterogeneous sources to unified metabolite entities (Lu, 2025). Identifier mapping errors cascade through research quality: incorrect links lead to spurious discoveries, failed recognition causes missed connections, and inconsistent usage prevents reproducible science (Cajka and Fiehn, 2024; Powers et al., 2022).

Database inconsistency compounds these challenges. Analysis of 11 biochemical databases revealed that while HMDB maintains high consistency (only 1.7% of names linked to multiple IDs), ChEBI and KEGG show 14.8% and 13.3% ambiguity respectively (Grapov et al., 2018). Inter-database mapping using metabolite names shows inconsistencies ranging from 0% to 81.2%, with similar results (0-83%) when mapping via reference identifiers (Grapov et al., 2018). Current solutions include manually curated mapping tables (e.g., MetabolitesID (Bioconductor, 2020)), web-based conversion tools like the Chemical Translation Service (Wohlgemuth et al., 2010), and APIs, but these face persistent challenges: no single resource provides comprehensive coverage, databases update frequently requiring continuous curation, automated tools struggle with isomers and context-dependent synonyms, and different tools employ incompatible interfaces (Swainston et al., 2014; Krettler et al., 2024). Large Language Models promise flexible, context-aware resolution leveraging both structured databases and unstructured literature. However, MetaBench demonstrates that current models exhibit catastrophic failure without retrieval augmentation, achieving less than 1% accuracy even for frontier architectures. This finding establishes identifier grounding as a critical bottleneck impeding the computational transformation of metabolomics toward integrated, AI-assisted discovery systems.