

Disentangling Ambiguity from Instability in Large Language Models: A Clinical Text-to-SQL Case Study

Angelo Ziletti*, Leonardo D'Ambrosi

Bayer AG
angelo.ziletti@bayer.com

Abstract

Deploying large language models for clinical Text-to-SQL requires distinguishing two qualitatively different causes of output diversity: (i) input ambiguity that should trigger clarification, and (ii) model instability that should trigger human review. We propose CLUES, a framework that models Text-to-SQL as a two-stage process (interpretations \rightarrow answers) and decomposes semantic uncertainty into an ambiguity score and an instability score. The instability score is computed via the Schur complement of a bipartite semantic graph matrix. Across AmbigQA/SituatedQA (gold interpretations) and a clinical Text-to-SQL benchmark (known interpretations), CLUES improves failure prediction over state-of-the-art Kernel Language Entropy. In deployment settings, it remains competitive while providing a diagnostic decomposition unavailable from a single score. The resulting uncertainty regimes map to targeted interventions - query refinement for ambiguity, model improvement for instability. The high-ambiguity/high-instability regime contains 51% of errors while covering 25% of queries, enabling efficient triage.

Keywords: Uncertainty Quantification, Text-to-SQL, Semantic Entropy, Clinical NLP, Electronic Health Records

1. Introduction

Text-to-SQL systems powered by Large Language Models (LLMs) promise to democratize access to Electronic Health Records (EHRs) and claims databases, enabling epidemiologists to query complex data using natural language (Lee et al., 2023; Ziletti and D'Ambrosi, 2024; Koretsky et al., 2025). While retrieval-augmented approaches (RAG) have shown promising results (Ziletti and D'Ambrosi, 2024), deployment faces a critical barrier: the risk of syntactically correct but semantically erroneous queries that return plausible yet incorrect results.

A key challenge is *latent ambiguity*: ambiguity that is not evident to the user but significantly affects results (Saparina and Lapata, 2024, 2025). Consider "How many patients over 18 have atopic dermatitis?": does "over 18" refer to age at diagnosis or current age? First diagnosis or any? Each interpretation yields a different SQL query and potentially different results. Left unresolved, such ambiguities undermine reproducibility of epidemiological research (Zozus et al., 2016; Denaxas et al., 2017).

Latent ambiguity also complicates uncertainty estimation. When a model produces diverse outputs, is this diversity a legitimate reflection of query ambiguity, or a sign of model instability? Existing methods provide a single score that conflates two fundamentally distinct sources: uncertainty arising from inherent ambiguity in the input (akin to aleatoric uncertainty), and uncertainty stemming from the model's internal instability or lack of knowledge (akin to epistemic uncertainty). This distinction enables targeted diagnosis of system failures for model improvement. For a deployed system, it is also operationally critical: high ambiguity should

trigger a clarification dialogue, while high instability should flag the query for human review.

Contributions. (1) CLUES, a framework that decomposes semantic uncertainty into ambiguity (H_I) and conditional instability (H_{RI}) via the Schur complement of a bipartite similarity matrix, enabling different interventions (clarification vs. review) and applicable to any black-box LLM. (2) An interpretation-generation procedure for epidemiological questions and a clinical Text-to-SQL dataset with multiple interpretations.¹ (3) Empirical evidence across open-domain QA and clinical Text-to-SQL that CLUES improves failure prediction and yields regime-based error concentration useful for targeted routing.

Paper Organization. Sec. 3 formalizes our two-stage generative framework and introduces the Schur complement construction. We validate the decomposition across three settings of increasing complexity: open-domain QA with gold interpretations (Sec. 4), a real-world clinical Text-to-SQL dataset (Sec. 5), and production deployment with on-the-fly interpretation generation (Sec. 6).

2. Related Work

Semantic entropy and kernel-based uncertainty. Semantic entropy clusters generations into equivalence classes for hallucination detection (Kuhn et al., 2023; Farquhar et al., 2024); Kernel Language Entropy (KLE) generalizes this by encoding pairwise semantic similarities. (Nikitin et al., 2024) Recent work derives kernel entropy via bias-variance-covariance decomposition, (Gruber and

¹<https://github.com/Bayer-Group/clues-clinical-nlp>

Buettner, 2024) combines semantic signals with token-level uncertainty, (Raghuvanshi et al., 2025) and exploits geometric measures such as semantic density and volume. (Li et al., 2025b,a; Grewal et al., 2024) However, these methods operate on a single pool of outputs and do not model multi-stage generative processes or distinguish uncertainty sources. (Nikitin et al., 2024; Gruber and Buettner, 2024)

Decomposing aleatoric and epistemic uncertainty. This distinction is fundamental in probabilistic machine learning and critical for deployment in safety-critical domains. (Kendall and Gal, 2017; Hüllermeier and Waegeman, 2021) Recent LLM work trains meta-models or uses temperature sensitivity to identify epistemic uncertainty, but these approaches require additional supervision or pipeline changes. (Ahdritz et al., 2024; Foodei et al., 2025; Song et al., 2025) In the EHR-QA domain, Kim et al. (2022) decompose uncertainty to detect ambiguous questions, but operate at the token level without modeling multi-stage generation.

Text-to-SQL and clinical uncertainty. Text-to-SQL has evolved from cross-domain benchmarks (Yu et al., 2018) to clinical applications for EHR access and clinical trial recruitment (Ziletti and D’Ambrosi, 2025; Ghosh et al., 2025; Deng et al., 2022). Benchmarks like EHRSQL (Lee et al., 2023) and BiomedSQL (Koretsky et al., 2025) evaluate clinical Text-to-SQL, while multi-turn approaches improve robustness (Ryu et al., 2024). Existing systems prioritize execution accuracy but provide limited uncertainty quantification (Kim et al., 2022; Jo et al., 2024; Ziletti and D’Ambrosi, 2024). Reliability work focuses on unanswerable detection, ensemble methods, or token entropy thresholds (Jo et al., 2024; Kim et al., 2024), but does not decompose uncertainty by source or model semantic equivalence.

3. Methodology: A Decomposition Framework for Semantic Uncertainty

Our methodology extends the geometric view of uncertainty from KLE to a multi-stage generative process. We construct a bipartite semantic graph representing the entire system and use tools from linear algebra and spectral graph theory to decompose its structural complexity.

3.1. Interpretation-Augmented Generation

We formalize the Text-to-SQL system as a two-stage generative process. Given an initial natural language query q :

1. **Interpretation Stage:** The system first generates a set of N distinct semantic interpretations, $\mathcal{I} = \{I_1, I_2, \dots, I_N\}$. Each interpretation I_n is a natural language reformulation representing one plausible reading of the original query q . This stage explicitly surfaces the ambiguity (or lack thereof) inherent in the input.
2. **Generation Stage:** For each interpretation $I_n \in \mathcal{I}$, the system generates M candidate SQL queries, executes them against the database, and verbalizes the retrieved data into natural language answers. This yields a set of answers $\mathcal{R}_n = \{R_{n,1}, R_{n,2}, \dots, R_{n,M}\}$ per interpretation. The complete answer set is $\mathcal{R} = \{R_{1,1}, \dots, R_{N,M}\}$, containing $N \times M$ total answers.

This one-to-many ($N \rightarrow N \times M$) structure allows us to distinguish diversity *across* interpretations (ambiguity) from instability *within* the generation process for a single interpretation.

3.2. Constructing the Bipartite Semantic Graph

We define a similarity function $k(\cdot, \cdot) \in [0, 1]$ measuring semantic similarity between text strings. The full system matrix \mathbf{W} serves as the graph’s weighted adjacency matrix and has a natural bipartite block structure:

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}_{II} & \mathbf{W}_{IR} \\ \mathbf{W}_{RI} & \mathbf{W}_{RR} \end{pmatrix} \in \mathbb{R}^{(N+NM) \times (N+NM)} \quad (1)$$

where:

- $\mathbf{W}_{II} \in \mathbb{R}^{N \times N}$: The *interpretation similarity matrix*, with $(\mathbf{W}_{II})_{ij} = k(I_i, I_j)$. This block encodes the semantic relationships among interpretations.
- $\mathbf{W}_{RR} \in \mathbb{R}^{NM \times NM}$: The *answer similarity matrix* (flattening $(n, m) \mapsto j$), with $(\mathbf{W}_{RR})_{ij} = k(R_i, R_j)$. This block encodes the semantic relationships among all generated answers.
- $\mathbf{W}_{IR} \in \mathbb{R}^{N \times NM}$ and $\mathbf{W}_{RI} = \mathbf{W}_{IR}^\top$: The *assignment matrices*, with $(\mathbf{W}_{IR})_{ij} = 1$ if answer R_j was generated from interpretation I_i , and 0 otherwise.

The diagonal blocks \mathbf{W}_{II} and \mathbf{W}_{RR} encode graded semantic similarity, interpreted as weighted edge connectivity in the graph. The off-diagonal blocks \mathbf{W}_{IR} and \mathbf{W}_{RI} encode the generative provenance between interpretations and their answers. We represent this provenance as binary connectivity, since each answer traces to exactly one interpretation. This mixed representation

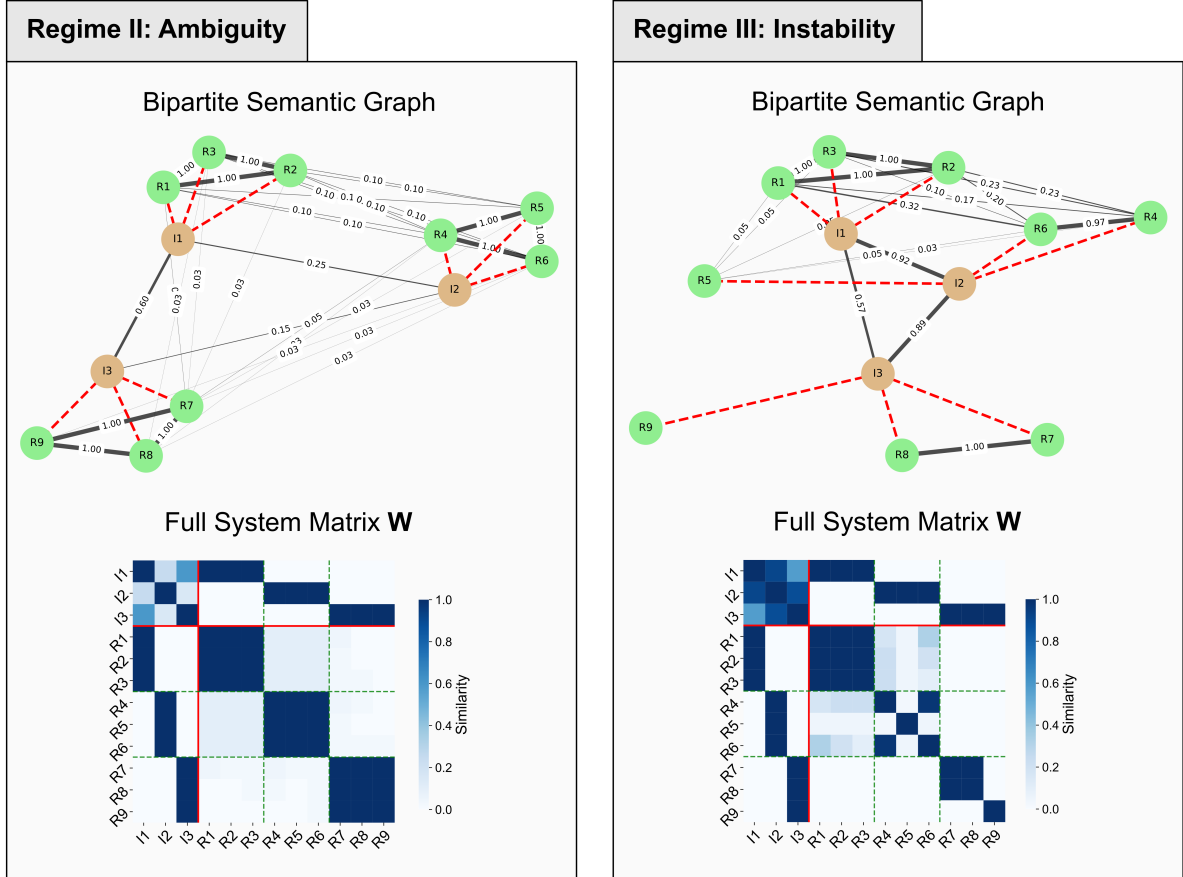


Figure 1: Examples of uncertainty regimes in CLUES. (Left) Regime II, Ambiguity: Interpretations (brown) are semantically distinct (low \mathbf{W}_{II} off-diagonal), but results (green) cluster tightly within each interpretation (high intra-cluster \mathbf{W}_{RR}), yielding high H_I , low $H_{R|I}$. (Right) Regime III, Instability: Interpretations are similar (high \mathbf{W}_{II} off-diagonal), yet results vary substantially within interpretations, yielding low H_I , high $H_{R|I}$. Red dashed edges: interpretation-result assignments (\mathbf{W}_{IR}). Bottom: corresponding system matrices \mathbf{W} .

is consistent with the graph-theoretic framework: edge weights encode connection strength, whether derived from continuous similarity scores or binary structural links. The KLE framework (Nikitin et al., 2024) constructs graphs from natural-language-inference-based similarity scores; our extension introduces a second node type (answers) connected to the first (interpretations) via known generative links rather than inferred similarity.

Prompt-Based Similarity. The similarity function $k(\cdot, \cdot)$ is implemented via an LLM with a task-specific prompt that defines the relevant notion of equivalence. For Text-to-SQL, the prompt assesses whether two interpretations would yield logically equivalent SQL queries; for QA benchmarks, whether they would yield the same factual answer. This ensures that \mathbf{W}_{II} and \mathbf{W}_{RR} encode task-relevant semantics rather than generic lexical similarity.

Heat Kernel Regularization. Following KLE (Nikitin et al., 2024), we apply a heat

diffusion kernel $\mathbf{K}_\tau = e^{-\tau\mathbf{L}}$, where $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is the graph Laplacian and \mathbf{D} is the degree matrix. This transforms raw pairwise similarities into a smoothed representation where the hyperparameter τ controls granularity: small τ preserves fine-grained distinctions, while large τ diffuses similarity across the graph, merging nearby nodes into coarse clusters. To select τ without a validation set, we calibrate against an idealized baseline: \mathbf{W}_{RR} of perfectly identical items (all ones) should yield near-zero entropy (<0.001 bits). This yields $\tau = 10$ for our experiments. Sensitivity analysis (Table 4) confirms robustness across $\tau \in [2, 20]$.

3.3. Decomposing Uncertainty with the Schur Complement

Our central goal is to decompose the total uncertainty of the system into interpretable components: *input ambiguity* and *system instability*. The result-

ing scores H_I and $H_{R|I}$ are not intended to replace aggregate measures like H_R for failure prediction; rather, they provide complementary information about the *source* of uncertainty, enabling differential interventions (input refinement vs. generation review) that a single score cannot support.

3.3.1. Baseline: Result Entropy H_R

The state-of-the-art KLE approach (Nikitin et al., 2024) computes entropy over generated outputs without modeling their provenance. In our framework, this yields H_R from \mathbf{W}_{RR} alone. High H_R indicates diverse outputs but does not distinguish the source of uncertainty: input ambiguity or system instability. This is the baseline we aim to improve upon.

3.3.2. Joint Uncertainty $H(R, I)$

We define total system uncertainty as the joint entropy over the full bipartite graph, computed from the complete similarity matrix \mathbf{W} (Eq. 1). This captures uncertainty across both interpretations and results, along with their generative relationships encoded in \mathbf{W}_{IR} .

3.3.3. Ambiguity Score H_I

We define the ambiguity score as the entropy over the set of interpretations, estimating ambiguity in the user’s query. We compute H_I by applying the KLE framework to \mathbf{W}_{II} :

$$\begin{aligned} \mathbf{K}_I &= e^{-\tau \mathbf{L}_I}, \quad \rho_I = \mathbf{K}_I / \text{Tr}(\mathbf{K}_I), \\ H_I &= -\text{Tr}(\rho_I \log \rho_I) \end{aligned} \quad (2)$$

3.3.4. Instability Score $H_{R|I}$

We define the instability score as the conditional entropy of results given interpretations: “Given a clear interpretation, how much does the answer still vary?”

The Naive Subtraction Approach. A natural approach is to estimate conditional entropy via subtraction: $H_{R|I} = H_{R,I} - H_I$. However, this relies on the Shannon entropy chain rule, which does not generally hold for von Neumann entropy on heat-kernel density matrices (Nikitin et al., 2024). The joint and marginal kernels have different sizes (see Eq. 1) and undergo independent diffusion processes, so $\tilde{H}_{R|I}$ is not guaranteed to be positive. Indeed, \tilde{H} yields negative values in 37–60% of cases (Tables 1 and 3).

The Schur Complement Approach. We leverage the block structure of \mathbf{W} (Eq. 1) to construct the conditional similarity structure directly via the Schur complement (Schur, 1917; Zhang, 2005):

$$\mathbf{S} = \mathbf{W}_{RR} - \mathbf{W}_{RI}(\mathbf{W}_{II} + \epsilon \mathbf{I})^{-1} \mathbf{W}_{IR} \quad (3)$$

where $\epsilon = 10^{-3}$ ensures invertibility. The Schur complement provides a principled way to “condition on” the interpretation structure. The projection term $\mathbf{W}_{RI} \mathbf{W}_{II}^{-1} \mathbf{W}_{IR}$ captures the portion of result similarity explained by interpretations. The residual \mathbf{S} retains only the similarity structure that input ambiguity cannot account for. This mirrors the Gaussian case, where the Schur complement yields conditional covariance $\text{Cov}(R|I)$ (Boyd and Vandenberghe, 2004). The entropy $H_{R|I}$ computed from \mathbf{S} therefore quantifies the system’s internal inconsistency. Since \mathbf{S} is not guaranteed to be positive semidefinite (PSD), we project onto the PSD cone via eigendecomposition, clipping negative eigenvalues to zero (Higham, 2002), then apply the KLE recipe (Eq. 2). Unlike the subtraction approach above, this guarantees $H_{R|I} \geq 0$. We term this decomposition framework **CLUES** (Conditional Language Uncertainty via Entropy and Schur). Crucially, CLUES distinguishes uncertainty from query ambiguity (potentially addressable via clarification) from system instability (requiring model-level intervention).

3.3.5. Uncertainty regimes and recommended interventions

We define four regimes by thresholding H_I (ambiguity) and $H_{R|I}$ (instability):

- Regime I: Confident (low H_I , low $H_{R|I}$): auto-answer.
- Regime II: Ambiguity (high H_I , low $H_{R|I}$): input refinement (clarification).
- Regime III: Instability (low H_I , high $H_{R|I}$): generation review.
- Regime IV: Compound (high/high): clarification + generation review.

In experiments we use median thresholds unless otherwise stated. We evaluate both failure prediction and regime-based routing in Sec. 4.2, 5.3, and 6.3.

4. Experiments on Open Answer Datasets

We first validate our decomposition on two open-domain QA benchmarks: AmbigQA (Min et al., 2020), where questions admit multiple valid interpretations due to semantic ambiguity, and SituatedQA (Zhang and Choi, 2021), where answers vary based on temporal or geographical context. Both datasets provide gold-standard interpretations paired with corresponding answers, allowing us to test whether $H_{R|I}$ provides a more accurate failure signal than H_R .

Model	AmbigQA ($N_q = 300$)								SituatedQA ($N_q = 300$)							
	Acc.	AUROC			Regime Acc.			Δ	Acc.	AUROC			Regime Acc.			Δ
		H_R	$\tilde{H}_{R I}$	$H_{R I}$	A	B	Δ			H_R	$\tilde{H}_{R I}$	$H_{R I}$	A	B	Δ	
Qwen3	25.0	.476	.550	.506	29.8	23.7	+6.2	20.7	.494	.521	.561	32.8	16.3	+16.5		
GPT-OSS	27.0	.641	.524	.702	53.8	8.1	+45.7	25.7	.660	.634	.758	44.4	3.1	+41.3		
KimiK2	40.7	.663	.466	.770	76.1	16.3	+59.7	37.3	.631	.590	.752	64.3	13.8	+50.5		
Claude4.5S	35.7	.501	.475	.533	40.7	35.7	+5.0	30.7	.502	.569	.495	44.4	19.5	+24.9		
Gemini3Pro	47.3	.538	.480	.659	73.8	32.6	+41.2	44.6	.584	.594	.685	61.6	12.7	+48.9		
<i>Pooled</i>	<i>35.1</i>	<i>.550</i>	<i>.498</i>	<i>.627</i>	<i>56.6</i>	<i>21.1</i>	<i>+35.5</i>	<i>31.8</i>	<i>.567</i>	<i>.573</i>	<i>.641</i>	<i>46.8</i>	<i>13.3</i>	<i>+33.5</i>		

Table 1: Multi-Model Validation on Open-Domain QA Benchmarks. We evaluate $N_q = 300$ questions per dataset across 5 LLMs (pooled: $N_q \times 5 = 1500$ question–model instances). We compare H_R (baseline KLE result entropy), naive subtraction $\tilde{H}_{R|I}$, and Schur-complement conditional entropy $H_{R|I}$ (CLUES) for failure prediction (AUROC). Δ denotes the accuracy gap between Regime A (high H_R , low $H_{R|I}$) and Regime B (high H_R , high $H_{R|I}$), where high/low are defined by pooled median splits. The naive subtraction $\tilde{H}_{R|I}$ yields negative values in 53%/37% of cases (AmbigQA/SituatedQA).

4.1. Experimental Setup

Dataset and Protocol. From each dataset, we select the first $N_q = 300$ questions with at least two gold interpretations, capping at three interpretations per question to limit inference costs. Following the two-stage framework in Sec. 3, the interpretation stage is given by the dataset’s gold interpretations. For the generation stage, we sample $M = 3$ answers per interpretation, yielding $N \times M$ results per question ($N \in \{2, 3\}$). We construct the bipartite semantic graph \mathbf{W} and compute H_R , $\tilde{H}_{R|I}$, and $H_{R|I}$.

Large Language Models. For answer generation, we evaluate five frontier LLMs: GPT-OSS-120B (hereafter *GPT-OSS*) (OpenAI et al., 2025), Kimi K2 Thinking (*KimiK2*) (Team et al., 2025) and Qwen-3-VL-235B-A22B (*Qwen3*) (Yang et al., 2025) are open-weights; Gemini 3 Pro (*Gemini3Pro*) (Gemini Team, 2023) and Claude Sonnet 4.5 (*Claude4.5S*) (Anthropic AI, 2023) are proprietary. All models are sampled with temperature $T = 1$. Throughout this paper, we use Gemini 2.5 Flash ($T = 0$) for semantic similarity (\mathbf{W}_{II} , \mathbf{W}_{RR}) and answer correctness evaluation (unless otherwise noted).

Evaluation Metric. We employ a strict path consistency strategy: a generated answer is correct only if it matches the gold answer for the specific interpretation that produced it, not any valid answer. Correctness is assessed via LLM-as-a-judge, prompting Gemini3Pro with the predicted and gold answers. A question is labeled as failure if the proportion of correct answers falls below a threshold η (we use $\eta = 0.8$). This criterion requires the system to respect the semantic constraints of each interpretation, penalizing both mode collapse (identical outputs across distinct interpretations) and context confusion (correct answer paired with wrong interpretation).

4.2. Results and Analysis

Results are presented in Table 1. Predicting failures on frontier LLMs is inherently challenging; the baseline H_R achieves AUROC scores of 0.47–0.66 depending on model and dataset. Our proposed $H_{R|I}$ consistently outperforms H_R , achieving pooled AUROC of 0.627 on AmbigQA (+0.077) and 0.641 on SituatedQA (+0.074). Per-model results show consistent improvements, with the largest gains when H_R itself carries predictive signal; for models where H_R performs near chance (Claude4.5S, Qwen3), all entropy measures show very limited discriminative power. The naive subtraction approach $\tilde{H}_{R|I}$ achieves near-chance AUROC, validating the need for the Schur complement construction (Eq.3).

The improvement over H_R stems from CLUES’s ability to account for input structure. By ignoring interpretations, H_R errs in both directions: *diversity inflation*, where legitimate variation across interpretations ($\mathbf{W}_{II} \approx \mathbf{I}$, $\mathbf{W}_{RR} \approx \mathbf{I}$) is flagged as uncertainty; and *mode collapse*, where identical outputs across distinct interpretations ($\mathbf{W}_{RR} \approx \mathbf{1}$) yield false confidence. The Schur complement detects both failure modes.

Regime Analysis. The four-regime framework (Sec. 3.3.5) partitions by H_I and $H_{R|I}$ to guide interventions. Here, we test whether $H_{R|I}$ provides discriminative value beyond H_R alone. We focus on high-uncertainty queries where H_R exceeds its median, then partition by median $H_{R|I}$ into Regime A (low $H_{R|I}$) and Regime B (high $H_{R|I}$). Despite identical H_R profiles, accuracy differs substantially: 56.6% vs. 21.1% for AmbigQA and 46.8% vs. 13.3% for SituatedQA. These differences are highly significant (Chi-square $p < 10^{-22}$), with odds ratios of 4.9 and 5.7: queries in Regime B are 5–6 \times more likely to fail. This pattern holds

Input Question	Disambiguated Question
How many patients > 17 yo have atopic dermatitis (all codes). Breakdown by code	Calculate the count of unique patients who have a first diagnosis of Atopic Dermatitis at any time in their patient history, where the patient's age at the time of this first diagnosis was greater than 17 years. Provide this count broken down by the specific concept code of Atopic Dermatitis.
How many patients with chronic kidney disease never took heparin before their chronic kidney disease diagnosis?	Count the number of unique patients who have a first occurrence of Chronic Kidney Disease and who have no record of Heparin administration at any time prior to their first Chronic Kidney Disease diagnosis.

Table 2: Disambiguation examples resolving temporal, demographic, and event ordering ambiguities. Bold text indicates specifications that were implicit or absent in the original question.

across all five LLMs without model-specific calibration.

5. Clinical Text-to-SQL with Known Interpretations

5.1. Resolving Ambiguity in Epidemiological Questions

Epidemiological questions in natural language frequently contain implicit ambiguities requiring explicit specification for accurate analysis (Table 2). While rule-based templates (e.g., OHDSI ATLAS (OHDSI Community, 2024)) lack flexibility and interactive clarification introduces user fatigue, our iterative disambiguation approach automatically infers the most plausible interpretation through successive refinement rounds, presenting users with a transparent final specification they can verify or edit before SQL generation. The disambiguation pipeline operates iteratively until convergence. Given an input question, each round prompts the LLM to:

1. Identify common epidemiological ambiguities including patient count semantics (unique patients vs. all records), temporal relationships (before/after, within timeframes), population definitions (inclusion/exclusion criteria), event ordering (first vs. any occurrence), and demographic specifications (age at diagnosis vs. current age).
2. Assign an ambiguity score $s \in [0, 1]$ and generate multiple candidate interpretations ordered by clinical plausibility.
3. Select the most plausible interpretation for the next round.

The prompt embeds domain conventions as defaults reflecting common practice in observational health research (e.g., standard age bands; defaulting to entire patient history when unspecified; unique patient counts). The process terminates

when ambiguity falls below a threshold ($s < 0.1$); in practice, within 1–3 iterations. We use Gemini 2.5 Flash with structured output schemas.

This iterative disambiguation offers three advantages: (1) transparency, as users can review and edit candidates before SQL generation; (2) educational value, as outputs serve as templates for well-formed questions; and (3) systematic application of epidemiological best practices without user-specified routine parameters.

In qualitative evaluation, epidemiologists confirmed that disambiguated outputs accurately reflect domain conventions and produce clinically appropriate disambiguations.

5.2. Experimental Setup

Source Dataset and Interpretation Construction. We build on EpiAskKB (Ziletti and D’Ambrosi, 2025), which contains question-SQL pairs for epidemiological research on EHR. Questions may be ambiguous, reflecting how epidemiologists naturally pose queries (Table 2). For each question, we construct a multi-interpretation setup following Sec. 4.1. From the question and gold SQL, we generate a disambiguated question reflecting the gold SQL, plus three alternative interpretations (Sec. 5.1). For each alternative, we generate SQL by prompting the LLM with the original question, alternative interpretation, and gold SQL to guide query structure. This constrained context minimizes structural errors. We use Gemini3Pro for all generation.

Generation and Execution Pipeline. For each question, we sample $M = 3$ SQL queries per interpretation, yielding $3 \times 3 = 9$ generated queries per model-temperature configuration. SQL generation, entity resolution, self-correction, and execution follow Ziletti and D’Ambrosi (2024). Queries are run on Optum’s de-identified Clinformatics® Data Mart Database, and retrieved results are verbalized into natural language answers, from which we compute H_R , $\tilde{H}_{R|I}$, and $H_{R|I}$ as in Sec. 3. Evaluation

follows Sec. 4.1 with $\eta = 0.6$.

5.3. Results

Results are presented in Table 3. $H_{R|I}$ outperforms H_R as a failure predictor in 17 of 20 model-temperature configurations, achieving pooled AUROC of 0.762 [95% CI: 0.737, 0.787] versus 0.600 [0.575, 0.627] for H_R . This difference is significant ($p < 10^{-10}$), demonstrating that conditioning on interpretation structure provides substantial discriminative value. The naive subtraction $\tilde{H}_{R|I}$ achieves only 0.461 AUROC with 60% negative values, validating the theoretical necessity of the Schur complement (Sec. 3.3).

Model (T)	Acc.	AUROC			
		H_R	$\tilde{H}_{R I}$	$H_{R I}$	Δ
<i>Qwen3</i>					
T=0.7	53.3	.605	.544	.764	+54.3
T=1.0	54.2	.617	.531	.748	+42.6
T=1.5	55.0	.618	.505	.818	+59.2
T=2.0	54.1	.630	.498	.786	+65.8
<i>GPT-OSS</i>					
T=0.7	71.6	.600	.416	.731	+31.2
T=1.0	75.5	.615	.410	.828	+34.3
T=1.5	71.6	.615	.384	.831	+34.2
T=2.0	71.6	.612	.391	.818	+35.0
<i>KimiK2</i>					
T=0.7	76.1	.632	.446	.815	+41.7
T=1.0	74.3	.548	.417	.817	+40.5
T=1.5	75.2	.635	.444	.791	+44.7
T=2.0	78.9	.599	.418	.853	+27.8
<i>Claude4.5S</i>					
T=0.7	78.9	.593	.395	.611	+0.0
T=1.0	77.6	.595	.466	.694	+10.3
T=1.5	78.0	.633	.400	.744	+30.8
T=2.0	78.9	.575	.426	.622	+16.8
<i>Gemini3Pro</i>					
T=0.7	93.3	.708	.623	.604	+3.8
T=1.0	93.3	.620	.532	.715	+20.6
T=1.5	88.8	.642	.578	.593	+5.9
T=2.0	88.8	.635	.542	.595	+10.7
<i>Pooled</i>	74.4	.600	.461	.762	+33.6

Table 3: Clinical Text-to-SQL benchmark results ($N = 2,180$; 109 questions \times 5 models \times 4 temperatures). $\Delta = \text{Regime A} - \text{Regime B}$ accuracy gap. $\tilde{H}_{R|I}$ produces negative values in 60% of cases.

Regime Analysis Confirms Discriminative Value. The accuracy gap Δ between Regime A (high H_R , low $H_{R|I}$) and Regime B (high H_R , high $H_{R|I}$), split at median values, is positive in 19 of 20 configurations (pooled: +33.6 pp). This replicates the open-domain QA pattern (Sec. 4.2): among

queries with equally high output diversity, those with low $H_{R|I}$ are substantially more likely to succeed.

Temperature Modulates Uncertainty Signal. For most models, moderate temperatures ($T = 1.0 - 1.5$) yield the strongest $H_{R|I}$ signal. Claude 4.5 illustrates this pattern: regime separation increases from $\Delta = 0$ at $T=0.7$ to $\Delta = +30.8$ at $T=1.5$ ($0 \rightarrow +30.8$), suggesting higher sampling diversity exposes latent instability. Gemini 3 Pro is the exception: its errors appear to stem from consistent but incorrect outputs rather than instability.

Sensitivity to Diffusion Parameter τ . Table 4 shows $H_{R|I}$ is robust for $\tau \geq 2$ (AUROC 0.75–0.76), while H_R degrades at high τ (0.74 \rightarrow 0.58) as diffusion blurs pairwise similarities toward uniformity. The Schur complement resists this effect: both raw similarities and the projection term smooth proportionally, preserving residual signal.

τ	1	2	3	5	10	15	20
H_I	.571	.574	.576	.578	.572	.569	.567
H_R	.669	.742	.738	.676	.601	.586	.582
$H_{R I}$.675	.746	.764	.765	.762	.761	.760

Table 4: Pooled AUROC sensitivity to heat kernel parameter τ . $H_{R|I}$ remains stable across $\tau \in [2, 20]$ while H_R degrades at high diffusion.

6. Clinical Text-to-SQL in Deployment Settings

6.1. Experimental Setup

Having established that $H_{R|I}$ outperforms H_R with known interpretations, we now test whether these gains extend to production deployment. This setting differs from previous experiments in three key ways.

Interpretations are generated, not given. Interpretations are produced on the fly via our disambiguation procedure (Sec. 5.1) and may imperfectly capture the true ambiguity structure, unlike the validated gold/silver labels used previously.

Uncertainty serves as a proxy signal. We evaluate final-answer correctness only, meaning $H_{R|I}$ computed over interpretations serves as a proxy for reliability rather than directly measuring per-interpretation accuracy.

Limited statistical power. Evaluation is coarser (one label per question rather than per interpretation), and models achieve high accuracy (Table 5), yielding sparse errors that challenge statistical estimation.

We use the same pipeline as Sec. 5.2, testing with and without disambiguation and RAG. RAG

follows the methodology in Ziletti and DAmbrosi (2024).

6.2. Results

Table 5 summarizes performance in the deployment setting. All models achieve 84–100% accuracy, higher than Sec. 5.3. This reflects the evaluation setup: strict path consistency requires correctness across multiple interpretations including less natural ones, while here we evaluate a single answer against the original query, which typically matches the model’s default interpretation.

		Qwen3	GPT-OSS	Kimik2	Claude4.5S	Gemini3Pro	Pooled
Dis	RAG						
×	×	85.3	93.6	99.1	95.4	100	94.7
✓	×	88.1	98.2	99.1	94.5	97.2	95.4
×	✓	90.8	99.1	97.2	97.2	100	96.9
✓	✓	94.5	92.7	94.5	96.3	100	95.6
Pooled		89.7	95.9	97.5	95.9	99.3	95.6

Table 5: Model Accuracy by Configuration. Dis. = disambiguated input. All values in %.

Uncertainty Separates Correct from Incorrect.

Despite sparse errors, median uncertainty values are systematically higher for incorrect predictions. Pooled across all configurations ($N = 2,180$), incorrect predictions show median H_I of 0.189 vs. 0.069 for correct, H_R of 0.442 vs. 0.301, and $H_{R|I}$ of 0.072 vs. 0.003.

RAG Stabilizes Output. RAG substantially reduces $H_{R|I}$ for correct predictions: GPT-OSS drops from 0.113 to 0.030, Qwen3 from 0.046 to 0.006, and Kimik2 from 0.062 to 0.025. However, RAG also reduces $H_{R|I}$ for incorrect predictions (pooled median: 0.104 \rightarrow 0.052), which may suppress the uncertainty signal that would otherwise flag errors (e.g., GPT-OSS: 0.296 \rightarrow 0.025). Claude4.5S and Gemini3Pro show minimal $H_{R|I}$ (0.002) with or without RAG, suggesting sufficient internalized domain knowledge.

AUROC and regime-based diagnostics. In this setting (Sec. 6.1), pooled AUROC is 0.687 [0.629, 0.740] for H_R versus 0.648 [0.600, 0.701] for $H_{R|I}$ (not significant, $p = 0.20$). Crucially, H_R cannot distinguish whether a high-uncertainty query needs clarification, model review, or both. CLUES is designed to *decompose* uncertainty rather than maximize single-score prediction, providing diagnostic resolution unavailable from H_R alone (Sec. 6.3).

6.3. Uncertainty Regime Analysis

Given sparse per-condition errors, we pool all data for regime analysis. We partition queries using pooled median thresholds for H_I and $H_{R|I}$, yielding a 2×2 structure (Table 6).

Regime	H_I	$H_{R I}$	C/I	Err
I: Confident	Low	Low	534/8	1.5%
II: Ambiguity	High	Low	530/18	3.3%
III: Instability	Low	High	529/21	3.8%
IV: Compound	High	High	492/48	8.9%
Total			2085/95	4.4%

Table 6: Uncertainty Regime Analysis. 2×2 decomposition by H_I and $H_{R|I}$ median thresholds. C/I = Correct/Incorrect.

Error Gradient and Routing Implications. Error rates increase from Regime I (1.5%) to Regime IV (8.9%). Each regime contains $\sim 25\%$ of queries, yet Regime IV concentrates 51% of all errors (48/95). A routing strategy sending only Regime IV to human review would examine one quarter of queries while catching half of all failures. Within-regime AUROC in Regime IV (≈ 0.52 for all metrics) indicates that continuous entropy values provide limited additional signal once the regime is determined. The decomposition’s value is further demonstrated among high-uncertainty queries where both H_I and H_R are above median ($N = 712$): these appear identical under conventional metrics, yet $H_{R|I}$ separates them into Regime A (4.3% error) and Regime B (9.2% error). This difference is significant (odds ratio OR = 2.24, $p < 0.02$) and holds across all five models using a single global threshold.

7. Summary

We introduced CLUES, a framework that decomposes semantic uncertainty into ambiguity (H_I) and conditional instability ($H_{R|I}$) via a Schur-complement construction over an interpretation–result bipartite graph. Across open-domain QA and clinical Text-to-SQL, $H_{R|I}$ is competitive with, and often improves upon, state-of-the-art Kernel Language Entropy (Nikitin et al., 2024) for failure prediction, while providing diagnostic resolution unavailable from a single score. Regime analysis based on $(H_I, H_{R|I})$ surfaces distinct failure patterns and maps them to targeted interventions: clarification for high ambiguity, human review for high instability. CLUES is model-agnostic and requires only output sampling; future work will explore adaptive routing and uncertainty propagation in agentic pipelines.

8. Limitations

Computational Cost. CLUES requires multiple LLM calls per query (interpretations \times samples; 9 in our setup). Calls can be parallelized, though cost remains a consideration for large-scale deployment.

Interpretation Quality. CLUES assumes generated interpretations meaningfully capture query ambiguity. If the disambiguation procedure produces trivial or redundant interpretations, H_I may underestimate true input ambiguity, and the Schur complement may not isolate instability effectively. We observed strong results with frontier LLMs, but interpretation quality likely degrades with weaker models.

Routing Efficiency in Sparse Error Regimes. In high-accuracy settings (95.6% in deployment), regime-based routing catches 51% of errors by reviewing 25% of queries, twice the efficiency of random selection, though not yet sufficient for fully automated triage. The fundamental challenge is sparse errors: with only 95 failures (\sim 4% error rate), fine-grained uncertainty signals have limited discriminative power. Applications with higher base error rates may benefit more substantially from regime-based routing.

9. Bibliographical References

- Gustaf Ahdrizt, Tian Qin, Nikhil Vyas, Boaz Barak, and Benjamin L. Edelman. 2024. Distinguishing the knowable from the unknowable with language models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org.
- Anthropic AI. 2023. Model card and evaluations for claude models. <https://www-files.anthropic.com/production/images/Model-Card-Claude-2.pdf>. Accessed: February 15, 2024.
- Stephen Boyd and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press, Cambridge, England.
- Spiros Denaxas, Kenan Direk, Arturo Gonzalez-Izquierdo, Maria Pikoula, Aylin Cakiroglu, Jason Moore, Harry Hemingway, and Liam Smeeth. 2017. [Methods for enhancing the reproducibility of biomedical research findings using electronic health records](#). *BioData Mining*, 10:31.
- Naihao Deng, Yulong Chen, and Yue Zhang. 2022. [Recent advances in text-to-SQL: A survey of what we have and what we expect](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2166–2187, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. [Detecting hallucinations in large language models using semantic entropy](#). *Nature*, 630(8017):625–630.
- Darius Foodeei, Simin Fan, and Martin Jaggi. 2025. [Semantic uncertainty in advanced decoding methods for llm generation](#). *arXiv preprint arXiv:2506.17296*.
- Gemini Team. 2023. [Gemini: A family of highly capable multimodal models](#). Technical report, Google. Accessed: February 15, 2024.
- Shrestha Ghosh, Moritz Schneider, Carina Reinicke, and Carsten Eickhoff. 2025. [A survey on LLM-assisted clinical trial recruitment](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 625–646, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Yashvir S. Grewal, Edwin V. Bonilla, and Thang D. Bui. 2024. [Improving uncertainty quantification in large language models via semantic embeddings](#). *arXiv preprint arXiv:2410.22685*.
- Sebastian G. Gruber and Florian Buettner. 2024. A bias-variance-covariance decomposition of kernel scores for generative models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org.
- Nicholas J. Higham. 2002. [Computing the nearest correlation matrix—a problem from finance](#). *IMA Journal of Numerical Analysis*, 22(3):329–343.
- Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506.
- Yongrae Jo, Seongyun Lee, Minju Seo, Sung Ju Hwang, and Moontae Lee. 2024. [LG AI research & KAIST at EHRSQL 2024: Self-training large language models with pseudo-labeled unanswerable questions for a reliable text-to-SQL system on EHRs](#). In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, pages 635–643, Mexico City, Mexico. Association for Computational Linguistics.
- Alex Kendall and Yarin Gal. 2017. [What uncertainties do we need in bayesian deep learning for computer vision?](#) *arXiv preprint arXiv:1703.04977*.

- Daeyoung Kim, Seongsu Bae, Seungho Kim, and Edward Choi. 2022. [Uncertainty-aware text-to-program for question answering on structured electronic health records](#). In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 138–151. PMLR.
- Sangryul Kim, Donghee Han, and Sehyun Kim. 2024. [ProbGate at EHRSQL 2024: Enhancing SQL query generation accuracy through probabilistic threshold filtering and error handling](#). In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, pages 687–696, Mexico City, Mexico. Association for Computational Linguistics.
- Mathew J Koretsky, Maya Willey, et al. 2025. BiomedSQL: Text-to-SQL for scientific reasoning on biomedical knowledge bases. *arXiv preprint arXiv:2505.20321*.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *International Conference on Learning Representations (ICLR)*.
- Gyubok Lee, Hyeonji Hwang, Seongsu Bae, Yeonsu Kwon, Woncheol Shin, Seongjun Yang, Minjoon Seo, Jong C Lee, and Edward Choi. 2023. EHRSQL: A practical text-to-SQL benchmark for electronic health records. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*.
- Xiaomin Li, Zhou Yu, Ziji Zhang, Yingying Zhuang, Swair Shah, Narayanan Sadagopan, and Anurag Beniwal. 2025a. [Semantic volume: Quantifying and detecting both external and internal uncertainty in llms](#). *arXiv preprint arXiv:2502.21239*.
- Zhaoye Li, Siyuan Shen, Wenjing Yang, Ruochun Jin, Huan Chen, Ligong Cao, and Jing Ren. 2025b. Enhancing uncertainty quantification in large language models through semantic graph density. In *Proceedings of the Forty-First Conference on Uncertainty in Artificial Intelligence*, UAI '25. JMLR.org.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. [AmbigQA: Answering ambiguous open-domain questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5783–5797, Online. Association for Computational Linguistics.
- Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. 2024. Kernel language entropy: fine-grained uncertainty quantification for llms from semantic similarities. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc.
- OHDSI Community. 2024. Atlas: Open source software for observational data analysis. <https://github.com/OHDSI/Atlas>. Accessed: 2026-01-28.
- OpenAI, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, Che Chang, Kai Chen, Mark Chen, Enoch Cheung, Aidan Clark, Dan Cook, Marat Dukhan, Casey Dvorak, Kevin Fives, Vlad Fomenko, Timur Garipov, Kristian Georgiev, Mia Glaese, Tarun Gogineni, Adam Goucher, Lukas Gross, Katia Gil Guzman, John Hallman, Jackie Hehir, Johannes Heidecke, Alec Helyar, Haitang Hu, Romain Huet, Jacob Huh, Saachi Jain, Zach Johnson, Chris Koch, Irina Kofman, Dominik Kundel, Jason Kwon, Volodymyr Kyrilov, Elaine Ya Le, Guillaume Leclerc, James Park Lennon, Scott Lessans, Mario Lezcano-Casado, Yuanzhi Li, Zhuohan Li, Ji Lin, Jordan Liss, Lily, Liu, Jiancheng Liu, Kevin Lu, Chris Lu, Zoran Martinovic, Lindsay McCallum, Josh McGrath, Scott McKinney, Aidan McLaughlin, Song Mei, Steve Mostovoy, Tong Mu, Gideon Myles, Alexander Neitz, Alex Nichol, Jakub Pachocki, Alex Paino, Dana Palmie, Ashley Pantuliano, Giambattista Parascandolo, Jongsoo Park, Leher Pathak, Carolina Paz, Ludovic Peran, Dmitry Pimenov, Michelle Pokrass, Elizabeth Proehl, Huida Qiu, Gaby Raila, Filippo Raso, Hongyu Ren, Kimmy Richardson, David Robinson, Bob Rotsted, Hadi Salman, Suvansh Sanjeev, Max Schwarzer, D. Sculley, Harshit Sikchi, Kendal Simon, Karan Singhal, Yang Song, Dane Stuckey, Zhiqing Sun, Philippe Tillet, Sam Toizer, Foivos Tsimpourlas, Nikhil Vyas, Eric Wallace, Xin Wang, Miles Wang, Olivia Watkins, Kevin Weil, Amy Wendling, Kevin Whinnery, Cedric Whitney, Hannah Wong, Lin Yang, Yu Yang, Michihiro Yasunaga, Kristen Ying, Wojciech Zaremba, Wenting Zhan, Cyril Zhang, Brian Zhang, Eddie Zhang, and Shengjia Zhao. 2025. [gpt-oss-120b and gpt-oss-20b model card](#). *arXiv preprint arXiv:2508.10925*.
- Shalini Raghuvanshi, Yashs Tiwari, and Abhishek Yadav. 2025. [Integrating token-level uncertainty, bidirectional nli, and semantic entropy for robust hallucination detection in large language models](#). In *2025 International Conference on Artificial Intelligence and Emerging Technologies (ICAJET)*, pages 1–6.

- Jaehee Ryu, Seonhee Cho, Gyubok Lee, and Edward Choi. 2024. [EHR-SeqSQL : A sequential text-to-SQL dataset for interactively exploring electronic health records](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16388–16407, Bangkok, Thailand. Association for Computational Linguistics.
- Irina Sapparina and Mirella Lapata. 2024. Ambrosia: a benchmark for parsing ambiguous questions into database queries. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24*, Red Hook, NY, USA. Curran Associates Inc.
- Irina Sapparina and Mirella Lapata. 2025. [Disambiguate first, parse later: Generating interpretations for ambiguity resolution in semantic parsing](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 16825–16839, Vienna, Austria. Association for Computational Linguistics.
- J Schur. 1917. Über potenzreihen, die im innern des einheitskreises beschränkt sind. *J. Reine Angew. Math.*, 1917(147):205–232.
- Haoyi Song, Ruihan Ji, Naichen Shi, Fan Lai, and Raed Al Kontar. 2025. [Inv-entropy: A fully probabilistic framework for uncertainty quantification in language models](#). *arXiv preprint arXiv:2506.09684*.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, Yichen Feng, Kelin Fu, Bofei Gao, Hongcheng Gao, Peizhong Gao, Tong Gao, Xinran Gu, Longyu Guan, Haiqing Guo, Jianhang Guo, Hao Hu, Xiaoru Hao, Tianhong He, Weiran He, Wenyang He, Chao Hong, Yangyang Hu, Zhenxing Hu, Weixiao Huang, Zhiqi Huang, Zihao Huang, Tao Jiang, Zhejun Jiang, Xinyi Jin, Yongsheng Kang, Guokun Lai, Cheng Li, Fang Li, Haoyang Li, Ming Li, Wentao Li, Yanhao Li, Yiwei Li, Zhaowei Li, Zheming Li, Hongzhan Lin, Xiaohan Lin, Zongyu Lin, Chengyin Liu, Chenyu Liu, Hongzhang Liu, Jingyuan Liu, Junqi Liu, Liang Liu, Shaowei Liu, T. Y. Liu, Tianwei Liu, Weizhou Liu, Yangyang Liu, Yibo Liu, Yiping Liu, Yue Liu, Zhengying Liu, Enzhe Lu, Lijun Lu, Shengling Ma, Xinyu Ma, Yingwei Ma, Shaoguang Mao, Jie Mei, Xin Men, Yibo Miao, Siyuan Pan, Yebo Peng, Ruoyu Qin, Bowen Qu, Zeyu Shang, Lidong Shi, Shengyuan Shi, Feifan Song, Jianlin Su, Zhengyuan Su, Xinjie Sun, Flood Sung, Heyi Tang, Jiawen Tao, Qifeng Teng, Chensi Wang, Dinglu Wang, Feng Wang, Haiming Wang, Jianzhou Wang, Jiaying Wang, Jinhong Wang, Shengjie Wang, Shuyi Wang, Yao Wang, Yejie Wang, Yiqin Wang, Yuxin Wang, Yuzhi Wang, Zhaoji Wang, Zhengtao Wang, Zhexu Wang, Chu Wei, Qianqian Wei, Wenhao Wu, Xingzhe Wu, Yuxin Wu, Chenjun Xiao, Xiaotong Xie, Weimin Xiong, Boyu Xu, Jing Xu, Jinjing Xu, L. H. Xu, Lin Xu, Suting Xu, Weixin Xu, Xinran Xu, Yangchuan Xu, Ziyao Xu, Junjie Yan, Yuzi Yan, Xiaofei Yang, Ying Yang, Zhen Yang, Zhilin Yang, Zonghan Yang, Haotian Yao, Xingcheng Yao, Wenjie Ye, Zhuorui Ye, Bohong Yin, Longhui Yu, Enming Yuan, Hongbang Yuan, Mengjie Yuan, Haobing Zhan, Dehao Zhang, Hao Zhang, Wanlu Zhang, Xiaobin Zhang, Yangkun Zhang, Yizhi Zhang, Yongting Zhang, Yu Zhang, Yutao Zhang, Yutong Zhang, Zheng Zhang, Haotian Zhao, Yikai Zhao, Huabin Zheng, Shaojie Zheng, Jianren Zhou, Xinyu Zhou, Zaida Zhou, Zhen Zhu, Weiyu Zhuang, and Xinxing Zu. 2025. [Kimi k2: Open agentic intelligence](#). *arXiv preprint arXiv:2507.20534*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chu-jie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#). *arXiv preprint arXiv:2505.09388*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Fuzhen Zhang, editor. 2005. *The schur complement and its applications*, 2005 edition. Numerical Methods and Algorithms. Springer, New York, NY.
- Michael Zhang and Eunsol Choi. 2021. [Situat-](#)

edQA: Incorporating extra-linguistic contexts into QA. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7371–7387, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Angelo Ziletti and Leonardo D'Ambrosi. 2024. Retrieval augmented text-to-SQL generation for epidemiological question answering using electronic health records. In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, pages 47–53, Mexico City, Mexico. Association for Computational Linguistics.

Angelo Ziletti and Leonardo D'Ambrosi. 2025. Generating patient cohorts from electronic health records using two-step retrieval-augmented text-to-sql generation. *arXiv preprint arXiv:2502.21107*.

Meredith N. Zozus, Rachel L. Richesson, Anita Walden, Jessie D. Tenenbaum, and W. E. Hammond. 2016. Research reproducibility in longitudinal multi-center studies using data from electronic health records. *AMIA Joint Summits on Translational Science Proceedings*, 2016:279–285.