

# Semantic Integrity Constraints: Declarative Guardrails for AI-Augmented Data Processing Systems

Alexander W. Lee  
Brown University  
alexander\_w\_lee@brown.edu

Justin Chan  
Brown University  
juchan@brown.edu

Michael Fu  
Brown University  
michael\_fu@brown.edu

Nicolas Kim  
Brown University  
nicolas\_kim@brown.edu

Akshay Mehta  
Brown University  
akshay\_mehta@brown.edu

Deepti Raghavan  
Brown University  
deeptir@brown.edu

Uğur Çetintemel  
Brown University  
ugur\_cetintemel@brown.edu

## ABSTRACT

*AI-augmented data processing systems (DPSs)* integrate large language models (LLMs) into query pipelines, allowing powerful semantic operations on structured and unstructured data. However, the reliability (a.k.a. trust) of these systems is fundamentally challenged by the potential for LLMs to produce errors, limiting their adoption in critical domains. To help address this reliability bottleneck, we introduce *semantic integrity constraints (SICs)*—a declarative abstraction for specifying and enforcing correctness conditions over LLM outputs in semantic queries. SICs generalize traditional database integrity constraints to semantic settings, supporting common types of constraints, such as grounding, soundness, and exclusion, with both reactive and proactive enforcement strategies.

We argue that SICs provide a foundation for building reliable and auditable AI-augmented data systems. Specifically, we present a system design for integrating SICs into query planning and runtime execution and discuss its realization in AI-augmented DPSs. To guide and evaluate our vision, we outline several design goals—covering criteria around expressiveness, runtime semantics, integration, performance, and enterprise-scale applicability—and discuss how our framework addresses each, along with open research challenges.

## PVLDB Reference Format:

Alexander W. Lee, Justin Chan, Michael Fu, Nicolas Kim, Akshay Mehta, Deepti Raghavan, and Uğur Çetintemel. Semantic Integrity Constraints: Declarative Guardrails for AI-Augmented Data Processing Systems. PVLDB, 18(11): 4073 - 4080, 2025.  
doi:10.14778/3749646.3749677

## 1 INTRODUCTION

Large language models (LLMs) have transformed the field of data management in the last couple of years, extending traditional systems with semantic processing capabilities over structured and unstructured data. In particular, several projects have integrated *semantic operators* into database and data flow systems by augmenting traditional operators with LLMs [2, 6, 14, 32, 36, 38, 44, 47, 49, 58, 60]. We broadly refer to these systems as *AI-augmented data processing systems (DPSs)*. While AI-augmented DPSs introduce novel ways to

process data, their use of LLMs lead to potential reliability issues that manifest into phenomena like hallucinations. These reliability issues hinder the adoption of AI-augmented DPSs in high-stakes sectors, such as healthcare, law, and finance. Without guardrails, it will be difficult to adopt these systems fully.

LLM assertions have actively been explored for offline evaluations [4, 25, 50, 51, 57, 66] and runtime guardrails [3, 16, 34, 35, 45, 53]. These assertions invoke custom code, ML classifiers, or LLM judges to check an LLM’s output with respect to various criteria. A related technique is constrained decoding, where constraints are applied during the LLM’s generation to ensure that the output adheres to a specified structure [9, 17, 21, 29, 30, 41, 64, 67]. While these approaches have proven their utility in enhancing LLM reliability, they are imperative and fragmented, limiting their usability.

To address these limitations, we propose *semantic integrity constraints (SICs)* as a declarative, unified abstraction that extends traditional database integrity constraints to support AI-augmented DPSs. Unlike traditional integrity constraints that prevent invalid modifications to the database, SICs protect against erroneous outputs from semantic operators. Similar to their traditional counterparts, SICs make defining constraints for predominant use cases easy and more advanced cases possible, aligning with recent work that argues for specifications on LLM-based system components [54].

This paper presents our vision for SICs and sketches a general framework for integrating them into AI-augmented DPSs. To evaluate the feasibility and impact of this vision, we define a set of design goals and success criteria: (1) *expressiveness*, to capture a broad range of common constraint types; (2) *executable semantics*, with well-defined runtime behavior; (3) *system integration*, via seamless support within data flow and query processing pipelines; (4) *performance-aware enforcement*, balancing reliability and cost; and (5) *enterprise-scale applicability*, including observability, constraint reuse, conflict detection, and auditability. These criteria, outlined in Table 1, were defined through our partnership with the Rhode Island Hospital. They guide our design and are reflected in the structure of the paper. Section 2 outlines the proposed system architecture. Section 3 introduces the SIC interface and categorizes SICs across key enforcement classes aligned with real-world LLM use cases. Section 4 explores enforcement strategies, spanning logical and physical planning, and outlines cost-aware mechanisms. Section 5 discusses observability components. Finally, Section 6 presents open challenges in large-scale constraint management.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 11 ISSN 2150-8097.  
doi:10.14778/3749646.3749677

**Table 1: Success criteria and design goals for SICs**

Success Criterion	Description
Expressiveness	SICs support a range of constraint types—including domain, grounding, and soundness—through a declarative interface (Section 3)
Executable Semantics	Each constraint has precise execution semantics—including what is checked, when it is enforced, and what actions are taken on failure—ensuring semantically consistent system behavior and predictable outcomes for users (Section 4)
System Integration	SICs integrate with query interfaces, optimizers, and execution engines in AI-augmented DPSs (Section 2)
Performance Awareness	Enforcement strategy is selected based on cost and reliability across reactive and proactive methods (Section 4.2)
Enterprise-Scale Applicability	SICs support observability, reuse, and auditability across workflows, with mechanisms for conflict detection and policy-driven enforcement at organizational scale (Sections 5 and 6)

## 2 SYSTEM ARCHITECTURE

We begin with an overview of the system architecture (Figure 1) for supporting SICs and defer a more detailed discussion of the core components to later sections. First, the user expresses their semantic query along with SICs via the *query interface* (Section 3.1). This interface forwards the query’s text to the *query parser*, which produces a corresponding logical query plan. Next, the *logical optimizer* (Section 4.1) applies logical rewrite rules to the initial plan, yielding an optimized logical plan. This plan is then passed to the *physical optimizer* (Section 4.2), which considers different physical implementations of the plan and selects the optimal one. The resulting optimized physical plan is then executed by the *execution engine*. To enable performance monitoring, an *observability stack* is integrated with the execution engine (Section 5). Finally, the *constraint store* recommends and analyzes SICs (Section 6).

## 3 CONSTRAINT SPECIFICATION

We suggest defining SICs declaratively, allowing users to focus on constraint logic while the optimizer determines the best execution method. We outline the main components of our proposed interface and detail classes of constraints relevant to semantic data processing. Additional surveys should be conducted to ensure usability.

### Listing 1: Semantic query with pipe SQL syntax and SICs.

```

1 FROM ehr_table
2 |> SET dob = p'canonicalize {dob} into YYYY-MM-DD'
3 |> ASSERT REGEXP_CONTAINS(dob,
   r'^\d{4}-(0[1-9]|1[0-2])-(0[1-9]|1[2]\d|3[01])$')
4 |> EXTEND DATE_PART('year', AGE(CURRENT_DATE,
   dob::DATE)) AS age_yrs
5 |> EXTEND EXTRACTIVE p'extract the patient's
   admission physical exam from the {ehr}' AS
   phys_exam STRING
6 |> EXTEND EXTRACTIVE p'extract the patient's
   admission lab results from the {ehr}' AS lab_res
   STRING
7 |> EXTEND EXTRACTIVE p'extract the patient's medical
   history from the {ehr}' AS med_hist STRING
8 |> EXTEND ABSTRACTIVE p'summarize {med_hist}' AS
   med_hist_sum STRING
9 |> WHERE p'the patient is likely to have sepsis based
   on their {age_yrs}, {phys_exam}, {lab_res}, and
   {med_hist}' AS sepsis_filter
10 |> ASSERT phys_exam GROUNDED AND lab_res GROUNDED
11 |> ASSERT med_hist_sum GROUNDED AND
   LENGTH(med_hist_sum) < 1000
12 |> ASSERT med_hist_sum EXCLUDES p'test results'
13   RETRY 1 CONTINUE ON FAIL
14 |> ASSERT sepsis_filter SOUND

```

### 3.1 Interface


To provide a unified interface for queries and SICs, users declare constraints in the query language. We present a relational interface that extends pipe SQL syntax [52] with semantic operators (e.g., map, aggregate, filter, join, top-*k*) and SICs; however, SICs are compatible with *any* declarative AI-augmented DPS.

**Declaring Constraints.** SICs specify *predicates* either on *generated tuple attributes* from semantic operators or on the *semantic operators* themselves. We make this distinction because while some operators (e.g., map, aggregate) generate attributes, other operators (e.g., filter, join, top-*k*) do not; yet, constraints should still be defined in either case. The syntax for declaring a constraint on an attribute involves a SQL expression starting with ASSERT, followed by a predicate involving the attribute’s name, the type of constraint (Section 3.2), and any necessary parameters. A constraint on an operator is similar. Users can also combine constraints via logical connectives (e.g., ASSERT <pred1> AND <pred2>).

**Retry Thresholds.** Users can define a retry threshold for each constraint. When a constraint associated with the operator (either on the operator or on an attribute generated by the operator) is violated, the DPS retries the operator until the output satisfies the constraint or the retry threshold is reached. To specify a retry threshold, users write RETRY <threshold> after the constraint.

**Failure Modes.** Failure modes (CONTINUE, IGNORE, or ABORT) are also associated with constraints and specify what the DPS should do when a failing output reaches the retry threshold. With CONTINUE, the DPS continues processing the query, allowing errors to propagate downstream. IGNORE specifies that the DPS should ignore invalid tuples in the processing (including dependent tuples, such as those that would be in the same aggregation window). Lastly, ABORT cancels the entire query. To specify a failure mode, the user writes <mode> ON FAIL after the retry threshold.

**End-To-End Example.** Listing 1 shows an example semantic query with various SICs. The query is declared over the `ehr_table` (line 1), a relation of electronic health records (EHRs). Each tuple in the relation corresponds to a patient, and we assume that it contains attributes for the patient’s `ehr` and `dob`. The query first canonicalizes the `dob` attribute’s value into YYYY-MM-DD (line 2) because the value could have been in another format. The expression involves a string prefixed with the letter “p”, which defines a *prompt string*. Prompt strings can specify natural language functions or predicates over attributes and are evaluated using LLMs. Line 4 computes the patient’s age using standard SQL functions. Lines 5-7 extract the patient’s physical exam, lab results, and medical history from the



**Figure 1: System architecture for SICs. Orange circles and turquoise diamonds represent semantic operators and SICs, respectively. Green squares indicate annotations that describe the physical implementation of each node.**

ehr, and line 8 summarizes the patient’s medical history. Lastly, the query filters tuples by selecting patients likely to have sepsis (line 9). Here, the AS keyword is also used to alias the filter operator so that it can be referenced later in the query (line 14).

The first constraint is declared on line 3, and the remaining constraints are specified on line 10 and onward. If the summarization step for med\_hist\_sum (line 8) fails the constraint on line 12, the DPS retries the operator once and then CONTINUEs to propagate the tuple even if it still fails the constraint again (line 13). Since no retry thresholds and failure modes are defined in the rest of the query, we assume that some default setting is used. We explain each constraint in more detail in the next section.

### 3.2 Constraint Classes

We categorize SICs into different classes (Table 2) that are designed to capture the predominant use cases for LLM constraints. We motivate each constraint class using concrete examples from the EHR query (among others) and briefly outline their implementation.

**Domain Constraints.** In many queries, users want generated attributes to belong to a specified *domain*. This includes use cases such as enforcing that a value conforms to a certain data type or structure, belongs to a set of values, or is less than a certain length. For example, one may check whether the predicted prescription dose from a semantic map is within a range of possible values. We extend traditional domain constraints to semantic data processing. In Listing 1, line 3 and the second conjunct on line 11 specify domain constraints, restricting the dob to the regex for YYYY-MM-DD and the summary’s length to less than 1000 characters, respectively. Type declarations (e.g., STRING) of generated attributes implicitly define domain constraints. This class of SICs comprises the majority of constraint use cases [34] and is enforced deterministically.

**Inclusion and Exclusion (IE) Constraints.** Users sometimes expect that an attribute generated by an operator *includes* or *excludes* certain terms, content, topics, or abstract features such as tones or styles. For example, given a semantic map that summarizes an EHR, users may want to exclude all the patient’s PII from the summary while including the names of attending physicians. IE constraints address this class of predicates. Line 12 defines an exclusion constraint on med\_hist\_sum, which excludes test results from the medical history summary. When an IE constraint is defined with a prompt string (as in the example), LLMs are used to check for a semantic match. Users can also specify IE constraints using strings, regexes, or even a collection of them (e.g., a set of toxic words to exclude from the output). In such cases, deterministic methods are used to evaluate matches.

**Grounding Constraints.** Semantic operators, such as maps or aggregates, can extract data or summarize from their inputs. However, since LLMs can hallucinate, users want to ensure that the generated data is *grounded* in source data. Grounding constraints enforce such invariants. For example, line 10 ensures that the extracted phys\_exam and lab\_res are *exactly contained* in the source EHR, while the first conjunct on line 11 asserts that the summarized med\_hist\_sum is *factually consistent* with the source EHR. In particular, the last grounding constraint implicitly defines a grounding constraint on med\_hist, since med\_hist\_sum is derived from med\_hist, which, in turn, is derived from ehr. Formally, if a grounding constraint is defined on an attribute, the DPS will recursively enforce grounding constraints on all ancestor attributes in its lineage. Thus, each constraint only needs to compare its operator’s input and output to ensure grounding in the source data. A grounding constraint is implemented in one of two ways, depending on whether the specified attribute is generated via an *extractive* (map) or *abstractive* (map or aggregate) operator. In the extractive case, the constraint is implemented deterministically and ensures that the operator’s output is precisely contained in its input. In the abstractive case, the constraint is implemented stochastically (e.g., with an LLM or a fact-checking model [56]) and ensures that the operator’s summarized output is factually consistent with the input. Users can annotate relevant operators with the EXTRACTIVE (lines 5-7) or ABSTRACTIVE (line 8) keywords to inform the DPS of their semantics; otherwise, the DPS infers these annotations by reasoning about the operator’s prompt string with an LLM.

**Soundness Constraints.** To check whether the reasoning of an operator is *logically sound*, users can declare soundness constraints. Although applicable to all operators, this constraint class is particularly useful for operators that do *not* generate attributes, such as filter, join, or top-*k*. To enforce soundness constraints, the DPS uses LLMs to check the operator’s reasoning, e.g., chain-of-thought (CoT) [62]. As an example, line 14 declares a soundness constraint on the sepsis filter. Assume that the query earlier retrieved relevant sepsis-related medical information (omitted for brevity). This retrieved information is considered to describe the symptoms of sepsis and the diagnostic rationale. Under the hood, the DPS prompts the filter’s LLM to return its CoT in addition to a boolean value that represents the selection result. The CoT first lists the premises, i.e., the patient’s medical condition and the retrieved medical information relevant to sepsis. The premises are followed by the reasoning steps that lead to the selection result. Then, the input and output of the filter’s LLM are given to an LLM judge (or several of them) to ensure that the premises are grounded in

**Table 2: Summary of SIC classes. The constraint classes are derived from the taxonomy introduced in prior work [34] and the use cases identified from our collaboration with the Rhode Island Hospital.**

Class	Description	Examples
Domain	Value is from a specified domain	Data types, structured output, multiple choices, numerical ranges, or length constraints
Inclusion/Exclusion	Value incl/excl terms, content, tones, styles, etc.	Inclusion of attending physician names or exclusion of PII from EHR summaries
Grounding	Value is grounded in the source values	Extracted test results or summaries are factual w.r.t. source EHRs
Soundness	LLM’s reasoning process is logically sound	Predictions of patient diagnoses are based on sound reasoning
Relevance	Value is relevant to the given task	Extracted histories from EHRs are medical (desired) rather than social (undesired)
Assertions	Arbitrary predicate holds in the data flow	All included URLs in generated research reports lead to valid websites

the input and that the reasoning steps are valid. The DPS can also use more sophisticated CoT prompting and verification techniques, which leverage structured natural language formats [31].

**Relevance Constraints.** These constraints validate if an attribute’s value is *relevant* for the given operator’s task; i.e., they verify whether the underlying LLM followed the instructions prompted by the user. Since relevance constraints are always appropriate, the DPS can enforce them on all generated attributes *by default*. We assume that this is the case in the example query. For instance, the DPS checks whether the medical history extracted in line 7 is in fact the patient’s medical history rather than something irrelevant, such as their social history. These constraints can be evaluated by passing the operator’s prompt, input, and output to an LLM judge.

**Assertions.** Like assertions in traditional integrity constraints, this class of constraints enables users to specify *arbitrary predicates* that must hold in the data flow. The DPS cannot offer many optimizations for assertions due to their black-box nature; however, such constraints give users the flexibility to express conditions that are not supported by the classes described previously.

## 4 CONSTRAINT ENFORCEMENT

Constraint enforcement consists of two phases: *detection* and *recovery*. First, a constraint *detects* any violations in the operator’s output. If a violation is detected, the operator attempts to *recover*. This distinction helps us reason about optimizing constraint enforcement, as discussed in this section.

### 4.1 Logical Optimization

The logical optimizer takes as input a logical query plan represented as a data flow graph, where each node is either an operator or a constraint. The goal of the logical optimizer is to minimize the cost of the plan by applying *rewrite rules* to semantic operators and constraints. Here, we refer to cost as any arbitrary cost metric, such as latency or monetary cost. To keep the presentation focused, we only discuss logical optimizations for constraints.

**Constraint Pushdown.** First, the optimizer moves all constraints to their earliest evaluation point in the data flow (i.e., immediately following their corresponding operators). This allows the DPS to detect and recover from constraint violations as soon as possible rather than unnecessarily processing erroneous outputs. The optimizer also appends the textual representation of each constraint to their corresponding operator’s LLM prompt, ensuring that the underlying model is aware of the constraints on its results.

**Constraint Reordering.** The optimizer then reorders the SICs for the same operator. Following prior work [1], the reordering is based on the cost and selectivity of each constraint, where selectivity is defined as the probability that a tuple satisfies the constraint’s predicate. Cost and selectivity estimates can be computed from sampled data. We discuss more details on estimation in Section 4.2.2.

### 4.2 Physical Optimization

Given an optimized logical plan, the physical optimizer considers different implementations of the plan and returns the optimal one. We first discuss approaches to constraint enforcement and then outline ideas for optimized plan selection.

**4.2.1 Enforcement Implementations.** Constraints are enforced *reactively* or *proactively*, based on whether enforcement occurs *after* or *during* response generation. While all constraints can be enforced reactively, only some can be enforced proactively.

**Reactive Enforcement.** After the LLM produces its entire response, there are different methods to detect a constraint violation. For example, a grounding constraint on an extractive operator can simply check if the operator’s output attribute is contained in the input. On the other hand, a grounding constraint on an abstractive operator can invoke a single LLM judge [66], a fact-checking model [56], a compound AI system [12], or anything in between. Similarly, there are different approaches for an operator to recover from a violation by retrying. The common approach used in current systems is *iterative refinement with feedback* [37, 49, 53]. Furthermore, the operator can retrieve additional semantically similar few-shot examples [10] of prior input-output pairs that satisfied the constraint. These examples can be stored in a cache built on a vector index. Other approaches include retrying with more powerful models, larger ensemble sizes, or higher thinking budgets [7, 20].

**Proactive Enforcement.** Proactively enforcing constraints can lower costs significantly. For example, *constrained decoding* [9, 17, 21, 29, 30, 41, 64, 67] pushes enforcement into the LLM’s decoding process, masking invalid tokens to ensure compliance by construction. Various libraries [17, 21, 41, 64, 67] support JSON Schema [42] adherence in LLM outputs by converting the schema into a grammar and then an automaton for token masking. However, constrained decoding is not limited to grammar-based approaches. For instance, using a suffix automaton is more efficient than regex to constrain an LLM to output a substring of its input [29], enabling proactive enforcement of grounding constraints for extractive operators. Yet, we are not aware of any existing library that offers this functionality.

An open challenge is to identify which other SICs can be enforced via constrained decoding. Moreover, developing a framework to apply efficiently various constraints on different output attributes also remains unresolved. Prior research [9] has seen some advancement but remains slower than modern grammar-based techniques [67]. Nonetheless, it provides a foundation for future investigation. Importantly, we suggest that the DPS should limit constrained decoding to simple constraints (seen by the LLM during training) to avoid sampling unlikely token distributions. This rationale also underlies why constraints are included in the LLM’s prompt before decoding.

Besides deterministic constrained decoding, the DPS can proactively enforce some stochastic constraints on partial outputs. For instance, the DPS can enforce a grounding constraint on an abstractive operator by invoking a fact-checking model [56] after each sentence is generated. While recent research [35] explores these ideas, they do not focus on efficient optimizations, such as *asynchronously* detecting violations in partial streamed outputs. This optimization is similar to previous work [48], except that the operator now needs to recover if its partial output violates a constraint. Upon a violation, the operator can backtrack to the decoding step right before the violation; though, alternative techniques like re-prompting with partial feedback are also worth considering.

**4.2.2 Optimized Plan Selection.** For a *fixed* set of constraints and their implementations, the physical optimizer can adapt existing techniques [32, 47] to select an optimal query plan defined by user-specified cost or reliability (a.k.a. quality) thresholds (omitted from the query interface for brevity). Specifically, rather than measuring the *reliability of an operator’s implementation* in terms of an oracle, reliability is instead defined by adherence to all of the operator’s constraints (after necessary retries). The challenge is selecting each constraint’s implementation in the first place.

As illustrated in Section 4.2.1, each constraint’s implementation has unique cost-reliability trade-offs, which allow users to specify additional thresholds to control them. While existing *cost* estimation approaches can be extended to account for each operator’s constraints, estimating the *reliability* of the constraints is a more difficult task. Formally, the *reliability of a constraint’s implementation* is defined as the *precision* and *recall* of detecting a violation. A deterministic constraint is completely reliable (i.e., 100% precision and recall) and requires no estimation, while a stochastic constraint is *not* completely reliable due to invoking models. Estimating the reliability of stochastic constraint implementations is challenging. First, such estimates require ground-truth labels from human annotators. Furthermore, reliability is a function of both the implementation and its input, and this input, in turn, depends on the upstream query plan and the underlying data. We present some initial ideas for constraint reliability estimation.

One strategy is to use *historical* data and workloads to estimate the precision and recall of different implementations for each constraint, providing these estimates to the physical optimizer for use at runtime. However, this approach neglects the dependencies mentioned earlier, leading to possibly noisy estimates. Since metrics are aggregated across diverse query contexts, they do not effectively reflect interactions with upstream operators. Furthermore, data and workload drift may make these estimates obsolete. It is also costly for individual organizations to annotate execution traces.

An alternative approach is to develop a *shared* foundation model for organizations, where given representations of (1) a constraint’s implementation, (2) the implementation of the query plan up until the constraint, and (3) the underlying data, the model predicts the precision and recall for the constraint’s implementation. Organizations can optionally fine-tune the model for their use cases. Recent work on foundation database models [61] are applicable here, with extensions to address the complexities of semantic processing.

A third approach is to use the model’s *confidence score* as a proxy for constraint reliability, allowing the optimizer to estimate the reliability (i.e., confidence) of a constraint’s implementation at runtime based on sampled data. Rather than defining precision and recall thresholds, users specify confidence thresholds instead. While many conventional classifiers inherently provide confidence scores, they can also be derived from LLMs by applying softmax to the “True” and “False” token logits. The difficulty lies in achieving calibrated confidence scores that accurately represent the implementation’s true correctness probability. Studies [26, 33] indicate that methods like temperature scaling [22] can enhance calibration in binary classification tasks. Though, further research is needed to obtain highly-calibrated constraints across varying domains.

## 5 OBSERVABILITY

After a user submits a query to the DPS, the observability stack enables the user to understand the performance of their query’s operators and constraints. The stack has three components. The *observability interface* allows the user to view relevant metrics (e.g., cost, reliability) and inspect query results. The *observability store* maintains metrics and outputs from each stage in the query. Lastly, the *labeling interface* enables human annotators to provide ground-truth labels to compute the precision and recall of each stochastic constraint. While methods such as LLM-as-a-judge offer a scalable approximation of human preferences [66] and have become widespread [3, 4, 16, 25, 37, 45, 49–51, 53, 57], they are not entirely reliable. As a result, labeling interfaces remain essential. We detail each component of the stack below.

**Observability Interface.** Given a user-specified query ID, the observability interface immediately shows metrics such as cost, operator reliability, and constraint selectivity. However, to observe the reliability of stochastic constraints, the user must submit a labeling request for their query. Once the labeling is complete, the interface displays the precision and recall of each constraint. The interface can also indicate whether a constraint is deterministic or stochastic; though, the query interface can expose this distinction through the EXPLAIN command as well. Moreover, the observability interface allows the user to inspect individual output tuples. A tuple is flagged if any upstream operator violated a constraint but CONTINUEd to propagate the error. The user can also view the tuple’s lineage for more fine-grained debugging.

**Observability Store.** The observability store is modeled as additional relations in the DPS. It maintains the lineage of each query, enabling systematic debugging and performance monitoring. To obtain human annotations, the DPS can draw concepts from crowd-sourced databases [19, 39, 43], such as crowd-sourced attributes and user interface generation; however, the DPS need not implement these ideas in their full generality. For example, suppose the

DPS stores a relation where each tuple represents an invocation of a constraint. The relation includes attributes for the constraint’s input and output (i.e., predicted label) among others (e.g., query ID). In particular, the relation also has a crowd-sourced attribute for the true label (initially null). When the user initiates a labeling request, the observability interface submits a query to the DPS, which computes (over a sample) the precision and recall of each constraint for the given query ID. The query populates each sampled true label by sending tasks to a crowd-sourcing platform.

**Labeling Interface.** A key aspect of crowd-sourced databases is their ability to create user interfaces automatically for crowd workers based on declarative queries. Here, the DPS produces a labeling interface for annotators, which displays a description of the constraint and its inputs. The annotator’s role is to determine if the inputs satisfy the constraint.

## 6 ENTERPRISE-WIDE CONSTRAINTS

The introduction of SICs into an enterprise environment presents both new opportunities and challenges. To effectively manage SICs within an organization, we propose the development of a *constraint store* and outline its key desired features.

**Recommendations.** Although declaring SICs is relatively simple, determining the applicable constraints is more difficult. This requires anticipating the potential failure cases for each semantic operator in the query. Recent systems have been developed to assist users in formulating constraints [50, 51], but these approaches typically generate suggestions based solely on the contents of the user’s query. In enterprise settings, constraint recommendations can be greatly enriched. Organizations typically enforce domain-specific rules that necessitate certain constraints for all pertinent queries. For instance, a hospital might require excluding PII in queries unless essential for analysis. Moreover, different teams within the same organization might use similar queries with shared constraints. Thus, the constraint store needs to capture and manage common constraints automatically organization-wide, facilitating their reuse and recommendation in related queries. We suggest using text embeddings [40, 46] and graph embeddings [28, 59, 65] to represent constraints and queries. This approach enables similarity searches to identify relevant constraints, which can then be refined by a reranker and LLM to match the user’s context.

**Conflicts.** Conflicts can arise when queries contain multiple SICs, such as from accumulating shared constraints across an organization’s workloads. For example, a query that requests a grounded list of the patient’s medical conditions with full details while also requiring the exclusion of PII may encounter conflicts when the relevant clinical information (e.g., diagnosis dates, provider names, hospital locations) are themselves considered PII, making it infeasible to satisfy the constraints without redaction or abstraction. Likewise, domain constraints over disjoint value sets can lead to inevitable failures. The constraint store allows for conflict analysis using static reasoning and LLM-based methods, issuing warnings. Users can resolve conflicts by adjusting priorities (e.g., making one constraint as soft) or relaxing scope as needed. This approach builds on early research in active databases [15, 23, 63], which suggested techniques for detecting rule conflicts and execution issues, such as non-termination, before execution [5].

## 7 RELATED WORK

Database integrity constraints were first described in Codd’s relational model [13] and later implemented in INGRES [24, 55] and System R [8, 11, 18]. Since then, constraints have been supported by all major databases. ML assertions were introduced to monitor and enhance model performance by applying arbitrary functions to model outputs to find potential errors [27]. Recent developments broaden assertions to LLMs by employing LLM judges and ML classifiers for offline evaluations [4, 25, 50, 51, 57, 66] and runtime guardrails [3, 16, 34, 45, 53]. Tools like SPADE [50] and EVALGEN [51] guide users by suggesting relevant assertions and their implementations (e.g., Python functions, LLM grader prompts). These tools complement SICs and can instead suggest declarative constraints to boost enforcement reliability and efficiency. Our vision for enterprise constraints also draws inspiration from these efforts. DSPy Assertions enable users to implement assertions in LLM pipelines as functions, which compile optimized prompts and iteratively refine outputs at runtime [53]. While SICs use iterative refinement, they do not currently support prompt optimization capabilities, although they can be added. A related area of work is controlled generation, which encompasses constrained decoding [9, 17, 21, 29, 30, 41, 64, 67] and model-based constraints [35]. To the best of our knowledge, there is no unified framework that offers declarative and optimized constraint enforcement.

Several recent systems have augmented traditional database operators with LLMs [2, 6, 14, 32, 36, 38, 44, 47, 49, 58, 60]. So far, DocETL is the only AI-augmented DPS to provide validation as a first class citizen [49]. They enable users to express simple Python statements and validation prompts for iteratively refining operator outputs. However, their approach to validation is largely imperative. LOTUS [44] and Palimpzest [32, 47] do not support constraints, but they allow users to specify quality thresholds relative to oracles.

## 8 CONCLUSIONS

SICs offer a unified, declarative foundation for improving the reliability of AI-augmented data processing. By generalizing traditional integrity constraints to the semantic behaviors of LLM-augmented operators, SICs bridge the gap between declarative query processing and the uncertainty of generative models. Their integration into the planning and execution layers of modern DPSs enables not only reliable enforcement but also cost-aware optimization, system-wide coordination, and organizational reuse. Through this framework, we outline a principled path toward building trustworthy AI-augmented systems capable of supporting mission-critical workflows. SICs are not only a guardrail mechanism, but also a catalyst for a new class of optimizations, abstractions, and research directions at the intersection of databases and LLMs.

## ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant Nos 2439559 and 2040433. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We thank Shu Chen, Duo Lu, Malte Schwarzkopf, Weili Shi, and the VectraFlow team for their valuable feedback.



## REFERENCES

- [1] Daniel J. Abadi, Don Carney, Uğur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. 2003. Aurora: a new model and architecture for data stream management. *The VLDB Journal* 12 (2003), 120–139. <https://doi.org/10.1007/s00778-003-0095-z>
- [2] Arun Agarwal and Renee Huang. 2025. Introducing Cortex AISQL: Reimagining SQL into AI Query Language for Multimodal Data. Retrieved 2025-07-08 from <https://www.snowflake.com/en/blog/ai-sql-query-language/>
- [3] Guardrails AI. 2023. Guardrails. Retrieved 2025-07-08 from <https://github.com/guardrails-ai/guardrails>
- [4] LangChain AI. 2023. LangChain. Retrieved 2025-07-08 from <https://github.com/langchain-ai/langchain>
- [5] Alexander Aiken, Jennifer Widom, and Joseph M Hellerstein. 1992. Behavior of database production rules: Termination, confluence, and observable determinism. In *Proceedings of the 1992 ACM SIGMOD international conference on Management of data*. 59–68.
- [6] Eric Anderson, Jonathan Fritz, Austin Lee, Bohou Li, Mark Lindblad, Henry Lindeman, Alex Meyer, Parth Parmar, Tanvi Ranade, Mehul A. Shah, et al. 2025. The Design of an LLM-powered Unstructured Analytics System. In *CIDR*.
- [7] Anthropic. 2025. Claude’s extended thinking. Retrieved 2025-07-08 from <https://www.anthropic.com/news/visible-extended-thinking>
- [8] M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade, and V. Watson. 1976. System R: relational approach to database management. 1, 2 (June 1976), 97–137. <https://doi.org/10.1145/320455.320457>
- [9] Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. 2023. Prompting Is Programming: A Query Language for Large Language Models. *Proc. ACM Program. Lang.* 7, PLDI, Article 186 (June 2023), 24 pages. <https://doi.org/10.1145/3591300>
- [10] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (NIPS ’20). Curran Associates Inc., Red Hook, NY, USA, Article 159, 25 pages.
- [11] D. D. Chamberlin, M. M. Astrahan, K. P. Eswaran, P. P. Griffiths, R. A. Lorie, J. W. Mehl, P. Reisner, and B. W. Wade. 1976. SEQUEL 2: A Unified Approach to Data Definition, Manipulation, and Control. *IBM Journal of Research and Development* 20, 6 (1976), 560–575. <https://doi.org/10.1147/rd.206.0560>
- [12] I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, and Pengfei Liu. 2023. FacTool: Factuality Detection in Generative AI—A Tool Augmented Framework for Multi-Task and Multi-Domain Scenarios. arXiv:2307.13528 [cs.CL] <https://arxiv.org/abs/2307.13528>
- [13] E. F. Codd. 1970. A relational model of data for large shared data banks. *Commun. ACM* 13, 6 (June 1970), 377–387. <https://doi.org/10.1145/362384.362685>
- [14] Hanjun Dai, Bethany Yixin Wang, Xingchen Wan, Bo Dai, Sherry Yang, Azade Nova, Pengcheng Yin, Phitchaya Mangpo Phothilimthana, Charles Sutton, and Dale Schuurmans. 2024. UQE: A Query Engine for Unstructured Databases. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 29807–29838. [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/34b3a40ec9752c1ae48fe85f8fe8dc-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/34b3a40ec9752c1ae48fe85f8fe8dc-Paper-Conference.pdf)
- [15] Umeshwar Dayal, Barbara Blaustein, Alex Buchmann, Upen Chakravarthy, Meichun Hsu, R Ledin, Dennis McCarthy, Arnon Rosenthal, Sunil Sarin, Michael J. Carey, et al. 1988. The HiPAC project: Combining active databases and timing constraints. *ACM Sigmod Record* 17, 1 (1988), 51–70.
- [16] Yi Dong, Ronghui Mu, Gaojie Jin, Yi Qi, Jinwei Hu, Xingyu Zhao, Jie Meng, Wenjie Ruan, and Xiaowei Huang. 2024. Building Guardrails for Large Language Models. arXiv:2402.01822 [cs.CL] <https://arxiv.org/abs/2402.01822>
- [17] Yixin Dong, Charlie F. Ruan, Yaxing Cai, Ruihang Lai, Ziyi Xu, Yilong Zhao, and Tianqi Chen. 2024. XGrammar: Flexible and Efficient Structured Generation Engine for Large Language Models. arXiv:2411.15100 [cs.CL] <https://arxiv.org/abs/2411.15100>
- [18] Kapali P. Eswaran and Donald D. Chamberlin. 1975. Functional specifications of a subsystem for data base integrity. In *Proceedings of the 1st International Conference on Very Large Data Bases* (Framingham, Massachusetts) (VLDB ’75). Association for Computing Machinery, New York, NY, USA, 48–68. <https://doi.org/10.1145/1282480.1282484>
- [19] Michael J. Franklin, Donald Kossmann, Tim Kraska, Sukriti Ramesh, and Reynold Xin. 2011. CrowdDB: answering queries with crowdsourcing. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data* (Athens, Greece) (SIGMOD ’11). Association for Computing Machinery, New York, NY, USA, 61–72. <https://doi.org/10.1145/1989323.1989331>
- [20] Google. 2025. Gemini 2.5 brings enhanced reasoning to enterprise use cases. Retrieved 2025-07-08 from <https://cloud.google.com/blog/products/ai-machine-learning/gemini-2-5-pro-flash-on-vertex-ai>
- [21] guidance. 2022. guidance-ai. Retrieved 2025-07-08 from <https://github.com/guidance-ai/guidance>
- [22] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) (ICML ’17). JMLR.org, 1321–1330.
- [23] Eric N. Hanson. 1996. The design and implementation of the Ariel active database rule system. *IEEE Transactions on Knowledge and Data Engineering* 8, 1 (1996), 157–172.
- [24] G. D. Held, M. R. Stonebraker, and E. Wong. 1975. INGRES: a relational data base system. In *Proceedings of the May 19-22, 1975, National Computer Conference and Exposition* (Anaheim, California) (AFIPS ’75). Association for Computing Machinery, New York, NY, USA, 409–416. <https://doi.org/10.1145/1499949.1500029>
- [25] Llama Index. 2023. Llama Index. Retrieved 2025-07-08 from [https://github.com/run-llama/llama\\_index](https://github.com/run-llama/llama_index)
- [26] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. 2022. Language Models (Mostly) Know What They Know. arXiv:2207.05221 [cs.CL] <https://arxiv.org/abs/2207.05221>
- [27] Daniel Kang, Deepti Raghavan, Peter Bailis, and Matei Zaharia. 2020. Model Assertions for Monitoring and Improving ML Models. In *Proceedings of Machine Learning and Systems*, Vol. 2. 481–496.
- [28] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SJU4ayYgl>
- [29] Terry Koo, Frederick Liu, and Luheng He. 2024. Automata-based constraints for language model decoding. In *First Conference on Language Modeling*. <https://openreview.net/forum?id=BDBdblmzy>
- [30] Michael Kuchnik, Virginia Smith, and George Amvrosiadis. 2023. Validating Large Language Models with ReLM. In *Proceedings of Machine Learning and Systems*, D. Song, M. Carbin, and T. Chen (Eds.), Vol. 5. Curran, 457–476. [https://proceedings.mlsys.org/paper\\_files/paper/2023/file/93c7d9da61ccb2a60ac047e92787c3ef-Paper-mlsys2023.pdf](https://proceedings.mlsys.org/paper_files/paper/2023/file/93c7d9da61ccb2a60ac047e92787c3ef-Paper-mlsys2023.pdf)
- [31] Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. 2023. Deductive Verification of Chain-of-Thought Reasoning. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 36407–36433. [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/72393bd47a35f5b3bee4c609e7bba733-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/72393bd47a35f5b3bee4c609e7bba733-Paper-Conference.pdf)
- [32] Chunwei Liu, Matthew Russo, Michael Cafarella, Lei Cao, Peter Baille Chen, Zui Chen, Michael Franklin, Tim Kraska, Samuel Madden, Rana Shahout, and Gerardo Vitagliano. 2025. Palimpsest: Optimizing AI-Powered Analytics with Declarative Query Processing. In *CIDR*.
- [33] Hongfu Liu, Hengguan Huang, Xiangming Gu, Hao Wang, and Ye Wang. 2025. On Calibration of LLM-based Guard Models for Reliable Content Moderation. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=wUbumOnd9N>
- [34] Michael Xieyang Liu, Frederick Liu, Alexander J. Fiannaca, Terry Koo, Lucas Dixon, Michael Terry, and Carrie J. Cai. 2024. “We Need Structured Output”: Towards User-centered Constraints on Large Language Model Output. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA ’24). Association for Computing Machinery, New York, NY, USA, Article 10, 9 pages. <https://doi.org/10.1145/3613905.3650756>
- [35] João Loula, Benjamin LeBrun, Li Du, Ben Lipkin, Clemente Pasti, Gabriel Grand, Tianyu Liu, Yahya Emara, Marjorie Freedman, Jason Eisner, Ryan Cotterell, Vikash Mansinghka, Alexander K. Lew, Tim Vieira, and Timothy J. O’Donnell. 2025. Syntactic and Semantic Control of Large Language Models via Sequential Monte Carlo. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=xoXn62FzD0>
- [36] Duo Lu, Siming Feng, Jonathan Zhou, Franco Solleza, Malte Schwarzkopf, and Uğur Çetintemel. 2025. VectraFlow: Integrating Vectors into Stream Processing. In *CIDR*.
- [37] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. SELF-REFINE: iterative refinement with self-feedback. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS ’23). Curran Associates Inc., Red Hook, NY, USA, Article 2019, 61 pages.

- [38] Samuel Madden, Michael Cafarella, Michael Franklin, and Tim Kraska. 2024. Databases Unbound: Querying All of the World’s Bytes with AI. *Proc. VLDB Endow.* 17, 12 (Aug. 2024), 4546–4554. <https://doi.org/10.14778/3685800.3685916>
- [39] Adam Marcus, Eugene Wu, David R. Karger, Samuel Madden, and Robert C. Miller. 2011. Crowdsourced Databases: Query Processing with People. In *CIDR*.
- [40] OpenAI. 2022. Introducing text and code embeddings. Retrieved 2025-07-08 from <https://openai.com/index/introducing-text-and-code-embeddings/>
- [41] OpenAI. 2024. Structured Outputs. Retrieved 2025-07-08 from <https://platform.openai.com/docs/guides/structured-outputs/>
- [42] JSON Schema Org. 2010. JSON Schema Spec. Retrieved 2025-07-08 from <https://github.com/json-schema-org/json-schema-spec>
- [43] Hyunjung Park, Hector Garcia-Molina, Richard Pang, Neoklis Polyzotis, Aditya Parameswaran, and Jennifer Widom. 2012. Deco: a system for declarative crowdsourcing. *Proc. VLDB Endow.* 5, 12 (Aug. 2012), 1990–1993. <https://doi.org/10.14778/2367502.2367555>
- [44] Liana Patel, Siddharth Jha, Parth Asawa, Melissa Pan, Carlos Guestrin, and Matei Zaharia. 2024. Semantic Operators: A Declarative Model for Rich, AI-based Analytics Over Text Data. arXiv:2407.11418 [cs.DB] <https://arxiv.org/abs/2407.11418>
- [45] Traian Rebedea, Razvan Dinu, Makesh Narsimhan Sreedhar, Christopher Parisien, and Jonathan Cohen. 2023. NeMo Guardrails: A Toolkit for Controllable and Safe LLM Applications with Programmable Rails. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Yansong Feng and Els Lefever (Eds.). Association for Computational Linguistics, Singapore, 431–445. <https://doi.org/10.18653/v1/2023.emnlp-demo.40>
- [46] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. arXiv:1908.10084 [cs.CL] <https://arxiv.org/abs/1908.10084>
- [47] Matthew Russo, Sivaprasad Sudhir, Gerardo Vitagliano, Chunwei Liu, Tim Kraska, Samuel Madden, and Michael Cafarella. 2025. Abacus: A Cost-Based Optimizer for Semantic Operator Systems. arXiv:2505.14661 [cs.DB] <https://arxiv.org/abs/2505.14661>
- [48] Keshav Santhanam, Deepti Raghavan, Muhammad Shahir Rahman, Thejas Venkatesh, Neha Kunjal, Pratiksha Thaker, Philip Levis, and Matei Zaharia. 2024. ALTO: An Efficient Network Orchestrator for Compound AI Systems. In *Proceedings of the 4th Workshop on Machine Learning and Systems* (Athens, Greece) (*EuroMLSys ’24*). Association for Computing Machinery, New York, NY, USA, 117–125. <https://doi.org/10.1145/3642970.3655844>
- [49] Shreya Shankar, Tristan Chambers, Tarak Shah, Aditya G. Parameswaran, and Eugene Wu. 2024. DocETL: Agentic Query Rewriting and Evaluation for Complex Document Processing. arXiv:2410.12189 [cs.DB] <https://arxiv.org/abs/2410.12189>
- [50] Shreya Shankar, Haotian Li, Parth Asawa, Madelon Hulsebos, Yiming Lin, J. D. Zamburescu-Pereira, Harrison Chase, Will Fu-Hinthorn, Aditya G. Parameswaran, and Eugene Wu. 2024. spade: Synthesizing Data Quality Assertions for Large Language Model Pipelines. *Proc. VLDB Endow.* 17, 12 (Aug. 2024), 4173–4186. <https://doi.org/10.14778/3685800.3685835>
- [51] Shreya Shankar, J.D. Zamburescu-Pereira, Bjoern Hartmann, Aditya Parameswaran, and Ian Arawjo. 2024. Who Validates the Validators? Aligning LLM-Assisted Evaluation of LLM Outputs with Human Preferences. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (*UIST ’24*). Association for Computing Machinery, New York, NY, USA, Article 131, 14 pages. <https://doi.org/10.1145/3654777.3676450>
- [52] Jeff Shute, Shannon Bales, Matthew Brown, Jean-Daniel Browne, Brandon Dolphin, Romit Kudtarkar, Andrey Litvinov, Jingchi Ma, John Morcos, Michael Shen, David Wilhite, Xi Wu, and Lulan Yu. 2024. SQL Has Problems. We Can Fix Them: Pipe Syntax In SQL. *Proc. VLDB Endow.* 17, 12 (Aug. 2024), 4051–4063. <https://doi.org/10.14778/3685800.3685826>
- [53] Arnab Singhvi, Manish Shetty, Shangyin Tan, Christopher Potts, Koushik Sen, Matei Zaharia, and Omar Khattab. 2024. DSPy Assertions: Computational Constraints for Self-Refining Language Model Pipelines. arXiv:2312.13382 [cs.CL] <https://arxiv.org/abs/2312.13382>
- [54] Ion Stoica, Matei Zaharia, Joseph Gonzalez, Ken Goldberg, Koushik Sen, Hao Zhang, Anastasios Angelopoulos, Shishir G. Patil, Lingjiao Chen, Wei-Lin Chiang, and Jared Q. Davis. 2024. Specifications: The missing link to making the development of LLM systems an engineering discipline. arXiv:2412.05299 [cs.SE] <https://arxiv.org/abs/2412.05299>
- [55] Michael Stonebraker. 1975. Implementation of integrity constraints and views by query modification. In *Proceedings of the 1975 ACM SIGMOD International Conference on Management of Data* (San Jose, California) (*SIGMOD ’75*). Association for Computing Machinery, New York, NY, USA, 65–78. <https://doi.org/10.1145/500080.500091>
- [56] Liyan Tang, Philippe Laban, and Greg Durrett. 2024. MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 8818–8847. <https://doi.org/10.18653/v1/2024.emnlp-main.499>
- [57] TruLens. 2021. TruLens. Retrieved 2025-07-08 from <https://www.trulens.org/>
- [58] Matthias Urban and Carsten Binnig. 2024. CAESURA: Language Models as Multi-Model Query Planners. In *CIDR*.
- [59] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rJXMpikCZ>
- [60] Jiayi Wang and Guoliang Li. 2025. AOP: Automated and Interactive LLM Pipeline Orchestration for Answering Complex Queries. In *CIDR*.
- [61] Johannes Wehrstein, Carsten Binnig, Fatma Özcan, Shobha Vasudevan, Yu Gan, and Yawen Wang. 2025. Towards Foundation Database Models. In *CIDR*.
- [62] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (*NIPS ’22*). Curran Associates Inc., Red Hook, NY, USA, Article 1800, 14 pages.
- [63] Jennifer Widom. 1996. The starburst active database rule system. *IEEE Transactions on Knowledge and Data Engineering* 8, 4 (1996), 583–595.
- [64] Brandon T. Willard and Rémi Louf. 2023. Efficient Guided Generation for Large Language Models. arXiv:2307.09702 [cs.CL] <https://arxiv.org/abs/2307.09702>
- [65] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=ryG6iA5Km>
- [66] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (*NIPS ’23*). Curran Associates Inc., Red Hook, NY, USA, Article 2020, 29 pages.
- [67] Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. 2024. SGLang: Efficient Execution of Structured Language Model Programs. arXiv:2312.07104 [cs.AI] <https://arxiv.org/abs/2312.07104>