



EVERGREEN: Efficient Claim Verification for Semantic Aggregates

Alexander W. Lee
Brown University and Snowflake Inc.
alexander_w_lee@brown.edu

Benjamin Han
Snowflake Inc.
benjamin.han@snowflake.com

Shayak Sen
Snowflake Inc.
shayak.sen@snowflake.com

Sam Yeom
Snowflake Inc.
sam.yeom@snowflake.com

Uğur Çetintemel
Brown University and Snowflake Inc.
ugur_cetintemel@brown.edu

Anupam Datta
Snowflake Inc.
anupam.datta@snowflake.com

ABSTRACT

With recent semantic query processing engines, semantic aggregation has become a primitive operator, enabling the reduction of a relation into a natural language aggregate using an LLM. However, the resulting semantic aggregate may contain claims that are not grounded in the underlying relation. Verifying such claims is challenging: they often involve quantifiers, groupings, and comparisons over relations that far exceed LLM context windows and require an expensive combination of semantic and symbolic processing.

We present EVERGREEN, a system that recasts claim verification as a semantic query processing task with tailored optimizations and provenance capture. EVERGREEN compiles each claim into a declarative semantic verification query that can execute on the same query engine used to produce the aggregate. To reduce cost, EVERGREEN avoids unnecessary LLM calls through verification-aware optimizations, including early stopping, relevance sorting, and estimation with confidence sequences, as well as general-purpose optimizations for semantic queries, such as operator fusion, similarity filtering, and prompt caching. Each verdict is accompanied by citations that identify a minimal set of tuples justifying the result, with semantics based on semiring provenance for first-order logic.

On a benchmark of production-inspired workloads over restaurant review and customer support datasets, EVERGREEN’s optimized configurations occupy the entire cost–quality Pareto frontier. With a strong LLM, EVERGREEN preserves verification quality at an F1 of 0.94 while reducing cost by 3.1× relative to unoptimized verification; with a substantially weaker LLM, it surpasses the strongest external baseline’s F1 (0.87 vs. 0.83) at 7.0× lower cost.

PVLDB Reference Format:

Alexander W. Lee, Benjamin Han, Shayak Sen, Sam Yeom, Uğur Çetintemel, and Anupam Datta. EVERGREEN: Efficient Claim Verification for Semantic Aggregates. PVLDB, 20(1): 16 - 30, 2026.
doi:10.14778/3845598.3845600

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/brown-db/evergreen>.

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. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 20, No. 1 ISSN 2150-8097.
doi:10.14778/3845598.3845600




Figure 1: F1 vs. mean cost per claim (log scale), averaged over three trials. We compare *base_rm* (LLM-as-a-judge), *rag_agent* (retrieval-augmented agent), *rlm* (recursive language model), *evg_unopt* (unoptimized EVERGREEN), and *evg_opt* (optimized EVERGREEN) across Claude and Qwen models. The dashed line connects the Pareto-optimal configurations. *evg_opt* exclusively occupies the Pareto frontier.

1 INTRODUCTION

Semantic query processing engines have recently emerged in both academia [4, 16, 18, 23, 47, 49, 51, 60, 65, 68, 75, 78] and industry [24, 32, 48, 82]. These systems feature *semantic operators* that augment database operators with LLMs. In several systems [4, 48, 60, 68], *semantic aggregation* is a primitive operator, reducing a relation of tuples with an LLM-based aggregation function specified by a prompt. For example, the prompt might summarize customer reviews or compare reviews across restaurants. When the relation fits in the LLM’s context window, the aggregation runs in a single LLM call; otherwise, it processes the relation in chunks and merges intermediate results. Yet, the resulting *semantic aggregate* may contain *claims* that are not *grounded* in the relation. The LLM might hallucinate that “the majority of reviews are positive” when in fact only a minority are, or that “no reviewers mention vegan options” when one does. The central question we investigate is: **how can we reliably and efficiently determine whether the claims embedded in a semantic aggregate are actually supported by the underlying relation?** As semantic aggregation gains adoption in production systems such as Snowflake’s Cortex AISQL [48], answering this question becomes essential.

Verifying claims in semantic aggregates is challenging for several reasons. First, the underlying relations are often large, easily exceeding LLM context windows; a relation of a few thousand reviews

can span hundreds of thousands of tokens. Second, from our observations of production workloads at Snowflake, claims that arise from such aggregations often involve quantifiers, groupings, and comparisons across groups (e.g., “all locations have multiple complaints”, “the top-ranked location is [A]”). Verifying these claims requires semantic reasoning about each tuple followed by symbolic reasoning over the resulting values, a task LLMs struggle to accomplish reliably. Third, semantically evaluating each tuple requires an expensive LLM invocation, making brute-force verification prohibitively costly. Finally, a boolean verdict alone is often insufficient: users expect citations to the underlying data that explain it.

Existing approaches to claim verification fall short for semantic aggregates. Naive LLM-as-a-judge methods [54, 72, 91] assess a claim against a provided source but assume that everything fits in the model’s context window, an assumption violated by large relations. Retrieval-augmented LLM approaches [41, 77, 80, 83] scale to larger corpora by retrieving relevant evidence, but they target claims verified by connecting dependent sources rather than quantified claims over independent tuples, and their natural language reasoning about quantifiers and groupings is costlier and less reliable than dedicated symbolic operations. Program-based LLM methods [34, 58, 87] can generate imperative programs that combine symbolic and semantic processing; however, they push the responsibility of optimization and citation capture onto the LLM, resulting in potentially suboptimal or error-prone implementations. In the database community, query-based verification systems [15, 35, 37, 40, 73, 74] translate claims into declarative queries over structured data, but cannot handle claims that require semantic evaluation. Meanwhile, semantic query processing engines are capable of expressing such verification logic. Yet, existing systems lack verification-aware optimizations and explanations. Most relevant to us is an early framework called BINDER [18], which naively processes all tuples with an LLM and was only evaluated on a claim verification benchmark limited to small structured tables [17].

We present EVERGREEN, a system that reliably and efficiently verifies claims from semantic aggregates by treating claim verification as a semantic query processing problem with principled optimizations and provenance capture. In particular, EVERGREEN compiles each claim into a declarative *semantic verification query* composed of standard relational operators and semantic operators, shifting the burden of optimization and citation capture away from the LLM and onto the query engine. Crucially, EVERGREEN can reuse the same query engine that produced the semantic aggregate, requiring no separate verification infrastructure. Its efficiency stems from two complementary classes of optimizations. First, EVERGREEN introduces *verification-aware optimizations*—including early stopping, relevance sorting, and estimation with anytime-valid confidence sequences [22, 79]—that exploit the structure of verification queries to minimize the amount of evidence that must be examined. These techniques are specific to verification queries and constitute our primary technical contributions. Second, EVERGREEN incorporates *general-purpose optimizations* for semantic queries, such as operator fusion, similarity filtering, and prompt caching. Finally, EVERGREEN accompanies each verdict with citations—a minimal explanation of why a claim holds or why not—formalized by applying and extending *semiring provenance* for first-order logic [31].

We evaluate EVERGREEN on a benchmark of 32 diverse claims over three Yelp restaurant review datasets [84] and three Twitter customer support datasets [10] based on production-inspired workloads. Figure 1 shows that across a range of LLMs, EVERGREEN’s optimized configurations form the entire cost–quality Pareto frontier. Compared to an unoptimized version of EVERGREEN with a strong LLM, its optimizations preserve verification quality at an F1 of 0.94 while reducing cost by 3.1×. The strongest external baseline [87] reaches an F1 of 0.83, which EVERGREEN surpasses at 0.87 with 7.0× lower cost using a substantially weaker LLM. This robustness reflects EVERGREEN’s declarative, neuro-symbolic design: an LLM specifies only the high-level verification query, while symbolic query processing handles optimization and execution, restricting LLMs to fine-grained semantic tasks where they are reliable.

In summary, we (1) identify several claim types that arise in semantic aggregates and map them to well-defined logical structures (Section 2); (2) show how such claims can be expressed as declarative semantic verification queries (Section 3); (3) introduce optimizations tailored to verification and integrate them with general-purpose optimizations for semantic queries (Section 4); (4) build on semiring provenance for first-order logic [31] to formalize citations as data provenance, deriving a minimal explanation for each verdict (Section 5); and (5) evaluate EVERGREEN on a benchmark reflecting real-world semantic aggregation workloads to demonstrate its dominance in verification quality and cost (Section 6).

2 CLAIMS OVER RELATIONS

Semantic aggregates implicitly encode a collection of *claims*, i.e., natural language assertions about the underlying relation. To verify these claims, EVERGREEN first makes them explicit via a *claim decomposition* process that invokes an LLM to break down each sentence in the semantic aggregate into a list of claims [80]. Since the resulting claims may contain vague references (e.g., pronouns, unknown entities), the process invokes an LLM again for each claim to resolve the references using the original response as context.

From our analysis of claims generated by production semantic aggregation workloads at Snowflake, we identify several common claim types that map to familiar logical structures. We first formalize our claim language and then describe each claim type.

Let $\mathcal{R}$ be a *relation*, which serves as the fixed, finite domain of our logic. $\mathcal{R}$ may be a filtered subset of the base relation relevant to the claim’s scope (e.g., reviews for a specific restaurant or reviews from vegetarian customers). That is, tuples not in $\mathcal{R}$ are not relevant to the claim. Let φ be a *quantifier-free formula* with a free variable ranging over tuples $t \in \mathcal{R}$ (i.e., a unary query), mapping each tuple to true if it satisfies the formula (i.e., $\varphi(t) = \top$) or false if it does not (i.e., $\varphi(t) = \perp$). φ may be built from symbolic predicates (e.g., whether the tuple’s age attribute is greater than 21) and semantic predicates (e.g., whether the tuple’s text attribute mentions vegan options) using conjunction, disjunction, and negation.

Existential claims use the existential quantifier to assert that *at least one* tuple satisfies φ , i.e., $\exists t \in \mathcal{R} : \varphi(t)$. For example, “Some reviewers enjoy the restaurant’s chicken salad” is existential.

Universal claims use the universal quantifier to assert that *all* tuples satisfy φ , i.e., $\forall t \in \mathcal{R} : \varphi(t)$. For example, “The reviews do not mention any vegan options” is universal.

Cardinal claims use cardinal quantifiers to assert the *number* of tuples that satisfy φ , such as $\exists_{\geq k} t \in \mathcal{R} : \varphi(t)$, where $k \in \{0, 1, \dots, |\mathcal{R}|\}$. For example, “At least a handful of vegetarian customers enjoyed the restaurant’s burgers” is a cardinal claim. Variants of cardinal quantifiers include $\exists_{>k}$, $\exists_{\leq k}$, $\exists_{<k}$, $\exists_{=k}$, and $\exists_{\neq k}$. Because $\mathcal{R}$ is fixed and finite, cardinal quantifiers generalize existential and universal quantifiers with $\exists_{\geq 1}$ and $\exists_{=|\mathcal{R}|}$, respectively.

Proportional claims use proportional quantifiers to assert the *proportion* of tuples that satisfy φ , such as $\exists_{\geq \rho} t \in \mathcal{R} : \varphi(t)$, where $\rho \in [0, 1]$. For example, “The restaurant has received majority positive reviews” is a proportional claim. Variants of proportional quantifiers include $\exists_{>\rho}$, $\exists_{\leq \rho}$, $\exists_{<\rho}$, $\exists_{=\rho}$, and $\exists_{\neq \rho}$. Proportional quantifiers reduce to cardinal quantifiers by scaling ρ by $|\mathcal{R}|$.

Ordinal claims assert that a group of tuples achieves a particular rank with respect to a per-group aggregate. Let $G_1, \dots, G_m \subseteq \mathcal{R}$ be groups of tuples induced by partitioning the tuples in $\mathcal{R}$ by one or more attribute values, and let $f(G_j, \varphi)$ denote a per-group aggregate (i.e., the count or proportion of tuples in G_j satisfying φ). Given the set $V(\varphi) := \{f(G_1, \varphi), \dots, f(G_m, \varphi)\}$ of distinct per-group aggregates, the rank of group G_i is

$$\text{rank}(G_i, \varphi) := |\{v \in V(\varphi) : v > f(G_i, \varphi)\}| + 1.$$

An ordinal claim asserts $\text{rank}(G_i, \varphi) = r$ for some target group G_i and rank r . For example, “The top-ranked McDonald’s location in terms of service has the Business ID [A]” is an ordinal claim.

Nested claims compose the prior claim types at multiple depths of grouping. Let $\mathcal{Q} := \{\exists, \forall, \exists_{\theta k}, \exists_{\theta \rho}\}$ denote the full set of quantifiers introduced earlier with $\theta \in \{\geq, >, \leq, <, =, \neq\}$. Given a partition of groups $\mathcal{G} := \{G_1, \dots, G_m\}$ as before, a nested claim of depth two specifies an outer quantifier $Q_{\text{out}} \in \mathcal{Q}$ over groups $G \in \mathcal{G}$ and an inner quantifier $Q_{\text{in}} \in \mathcal{Q}$ over tuples $t \in G$:

$$Q_{\text{out}} G \in \mathcal{G} : (Q_{\text{in}} t \in G : \varphi(t)).$$

For example, “All McDonald’s locations have multiple complaints about poor service quality” is $\forall G \in \mathcal{G} : (\exists_{\geq 2} t \in G : \varphi(t))$, where G is a group of tuples for a location and φ is the formula for whether a tuple complains about the restaurant’s poor service. Nesting can extend to arbitrary depths, though deeper nesting is rare in practice.

3 CLAIMS AS QUERIES

Based on the logical structure of claims, EVERGREEN uses an LLM to compile natural language claims into *semantic verification queries* composed of standard relational operators and semantic operators. Akin to traditional database query languages, these verification queries are *declarative*, enabling the LLM to focus on specifying *what* to verify rather than *how* to verify it. This shifts low-level responsibilities such as optimization and provenance capture away from the LLM and onto the query engine, improving reliability and efficiency. Moreover, verification requires no separate infrastructure: the same query engine that produced the semantic aggregate can also verify claims extracted from it.

Query Interface. EVERGREEN provides a DataFrame API in Python, where a DataFrame df represents a relation $\mathcal{R}$. Queries are composed by chaining operators parameterized by expressions and then optimized and executed using `collect()`. We detail the core operators and their relevant expressions, highlighting how they correspond to the logical structure of claims.

filter(predicate) selects tuples that satisfy the boolean expression predicate. In verification queries, filters are generally used to restrict $\mathcal{R}$ to tuples relevant to the claim. In addition to symbolic expressions for structured filtering, EVERGREEN features prompt expressions for semantic filtering: `prompt(prompt_str, return_type=bool)`. The prompt expression uses an LLM to evaluate a semantic predicate specified by a prompt string `prompt_str` that can reference tuple attributes by name.

```
df.filter(prompt("The {text} mentions the restaurant's service"))
```

map(expr) computes a new attribute for each tuple by evaluating `expr` to materialize attributes required by the claim’s formula φ that are not yet present in the relation. As with filters, `expr` may be a symbolic or prompt expression, though prompt expressions here can also return non-boolean types.

```
df.map(
    prompt("Identify the {text}'s sentiment", Sentiment).alias("sent")
)
```

aggregate(agg_exprs, group_by=()) aggregates over tuples by evaluating the sequence of aggregation function expressions `agg_exprs`, optionally grouped by a sequence of grouping expressions `group_by`. Aggregation functions are closely related to a claim’s quantifiers, and an aggregation function’s expression `expr` corresponds to the claim’s formula φ . `bool_or(expr)` corresponds to the existential quantifier, while `bool_and(expr)` corresponds to the universal quantifier. `count_if(expr)` evaluates to the number of tuples satisfying `expr`, and when the resulting aggregate is paired with a comparison expression in a subsequent operator, it corresponds to a cardinal quantifier. Similarly, `proportion(expr)` evaluates to the proportion of tuples satisfying `expr` and corresponds to a proportional quantifier when paired with a comparison expression. `expr` can simply be a reference `col(name)` to a boolean attribute or an expression over attributes that evaluates to a boolean.

```
df.aggregate(
    [proportion(col("sent").eq(Sentiment.POSITIVE)).alias("pos_prop")]
)
```

Grouping is used for ordinal and nested claims to compute group aggregates before subsequent ranking or outer aggregation steps.

with_rank(expr, descending=True) creates a dense rank attribute `col("rank")` for each tuple based on the expression `expr`, in descending order by default. This operator is used for ordinal claims, determining each group’s rank after per-group aggregation.

check(predicate) is the terminal operator of a verification query. It evaluates the boolean expression predicate on the (typically aggregated) result and adds it as a new boolean attribute representing the claim’s verdict. The verdict is annotated with provenance information that explains why the verdict is true or false, which we describe more in Section 5.

End-to-End Examples. Figure 2 shows the verification query for the nested claim from Section 2. The query first uses a semantic `map()` to evaluate φ on each tuple, labeling whether a review complains about the restaurant’s poor service. The first `aggregate()` groups tuples by location (`business_id`) and applies `count_if()` to count complaints per group for the inner cardinal quantifier $\exists_{\geq 2}$. The second `aggregate()` applies `bool_and()` to verify that every group’s complaint count meets the threshold, implementing the

```

df.map(
  prompt(
    "Identify whether the {text} complains about "
    "the restaurant's poor service", bool
  ).alias("service_complaint")
)
.aggregate(
  [count_if(col("service_complaint")).alias("complaint_count")],
  group_by=[col("business_id")]
)
.aggregate(
  [bool_and(col("complaint_count") >= 2).alias("all_mul_complaints")]
)
.check(col("all_mul_complaints"))
.collect()

```

Figure 2: Verification query for the nested claim “All McDonald’s locations have multiple complaints about poor service quality” ($\forall G \in \mathcal{G} : (\exists_{\geq 2} t \in G : \varphi(t))$).

```

df.filter(prompt("The {text} mentions the restaurant's service"))
.map(
  prompt(
    "Identify whether the {text} praises the restaurant's service",
    bool
  ).alias("praises_service")
)
.aggregate(
  [proportion(col("praises_service")).alias("service_praise_prop")],
  group_by=[col("business_id")]
)
.with_rank(col("service_praise_prop"))
.filter(col("business_id").eq("[A]"))
.check(col("rank").eq(1))
.collect()

```

Figure 3: Verification query for the ordinal claim “The top-ranked McDonald’s location in terms of service has the Business ID [A]” ($\text{rank}(G_i, \varphi) = 1$).

outer universal quantifier $\forall$. Finally, `check()` evaluates the verdict. Figure 3 shows a second example with an ordinal claim. The query additionally uses a semantic `filter()` to restrict to service-related reviews and `with_rank()` to determine each location’s rank.

4 QUERY OPTIMIZATIONS

The bottleneck of semantic verification queries lies in their semantic operators, which require invoking expensive LLMs. EVERGREEN introduces verification-aware optimizations and combines them with general-purpose optimizations for semantic queries to reduce cost while maintaining reliability.

4.1 Verification-Aware Optimizations

Early Stopping. Computing full aggregations over all tuples is not required for many verification queries, since the aggregates in our setting are *monotone* in the number of satisfying tuples. As such, EVERGREEN executes queries using the *iterator model*—where operators lazily pull tuples from their children—and stops aggregating early so that only necessary tuples are pulled through the semantic operators that appear before the aggregation. Table 1 summarizes the conditions under which each aggregate stops deterministically and probabilistically, and we now describe how they are realized.

`bool_or(expr)` stops aggregating after encountering a *witness* (i.e., a tuple that satisfies `expr`), while `bool_and(expr)` stops after encountering a *counterexample* (i.e., a tuple that does not satisfy `expr`). For `count_if(expr)` and `proportion(expr)`, the query optimizer extracts the relevant comparison expressions from subsequent operators (e.g., `check(col("pos_prop") > 0.5)`) and pushes them down to the corresponding aggregate operators as early stopping hints. Given these hints, an accumulator stops as soon as the comparison is settled: either its running count c already determines the outcome, or, when the total number n of input tuples is known, the result cannot change regardless of whether the remaining tuples satisfy `expr`. For an *ungrouped* aggregate, n is determined by a cardinality scan operator that counts the tuples prior to LLM evaluation. While scanning is considered expensive in traditional query processing, it is effectively free when compared to LLM evaluation. For a *grouped* aggregate, n is a per-group count computed during the sorting phase described next.

For grouped aggregates, EVERGREEN avoids using hash-based aggregates since they require pulling all tuples through the semantic operators, which defeats the purpose of early stopping. EVERGREEN instead uses streaming aggregates, processing groups sequentially over sorted, group-contiguous input. While streaming aggregates are classically only used when input tuples are *already* sorted, sorting (similar to scans) is cheap in our setting. The optimizer inserts a sort operator as low in the query plan as possible to ensure that tuples belonging to the same group are aggregated together while avoiding unnecessary semantic operator evaluations. Once a group’s aggregate is resolved, all operators below the aggregate skip the remaining tuples belonging to that group.

It is uncommon for traditional query engines to push comparison expressions into aggregates. This optimization is, however, reminiscent of prior research in ad-hoc sensor networks that pushes HAVING predicates into *grouped* monotonic aggregates to reduce transmission and storage of groups that do *not* satisfy the predicate [52]. There, as in our setting, per-tuple processing is relatively expensive. Yet, unlike this prior work, EVERGREEN pushes these comparisons down to *confirm* or *refute* them as early as possible for both *ungrouped* and *grouped* aggregates. We further amplify the benefits of this approach with the next two optimizations.

Relevance Sorting. Early stopping is most effective when relevant tuples—those that contribute to the aggregate’s verdict—appear early in the input stream. Since tuples are processed in arbitrary order by default, these relevant tuples may not surface until later. EVERGREEN encourages stopping by pre-sorting tuples by relevance.

At optimization time, the optimizer constructs a description of the aggregate to score each tuple’s relevance. It first traces the attribute referenced by the aggregate expression back to the semantic map that produced it, collecting the map’s prompt expression and any filter prompts that reference the same text attribute. These prompts provide context for what the aggregate is accumulating over (i.e., the scope of the claim’s relation $\mathcal{R}$ and the claim’s formula φ). The description also includes the aggregate expression itself and, when applicable, its comparison expression; together, they express the claim’s quantifier. The optimizer then prompts an LLM with this description to generate a *semantic search query*, a set of *inclusion keywords* likely to appear in relevant tuples, and a set of *exclusion keywords* that should not appear in them. At execution

Table 1: Deterministic and probabilistic stopping conditions to confirm/refute a claim. After observing s of n tuples with c satisfying expr , the final count is in $[c, c + (n - s)]$. $[L_s, U_s]$ is the proportion CI and ε the error tolerance. “—” marks a probabilistic condition subsumed by the deterministic one. $\exists_{>k}$ follows from $\exists_{\geq k+1}$. $\exists_{<k}$, $\exists_{\leq k}$, and $\exists_{\neq k}$ swap confirm/refute of $\exists_{\geq k}$, $\exists_{>k}$, and $\exists_{\leq k}$, respectively. Proportional quantifiers replace counts with proportions and k with ρ .

	Deterministic		Probabilistic	
	Confirm	Refute	Confirm	Refute
$\exists$	$c \geq 1$	$c + (n - s) < 1$	—	$nU_s < 1$
$\forall$	$c = n$	$c < s$	$L_s \geq 1 - \varepsilon$	—
$\exists_{\geq k}$	$c \geq k$	$c + (n - s) < k$	$nL_s \geq k$	$nU_s < k$
$\exists_{\leq k}$	$c = k \wedge s = n$	$c > k \vee c + (n - s) < k$	$[nL_s, nU_s] \subseteq [k(1 - \varepsilon), k(1 + \varepsilon)]$	$[nL_s, nU_s] \cap [k(1 - \varepsilon), k(1 + \varepsilon)] = \emptyset$

time, each tuple is scored using three ranking signals: cosine similarity between the tuple’s text embedding and the search query’s embedding, the number of inclusion keywords present, and the number of exclusion keywords absent. Each tuple’s text embedding is computed during ingestion and reused across queries. Signals are combined via Reciprocal Rank Fusion [19] (with the standard constant 60), and tuples are fed to the aggregate in descending relevance order. The relevance sort operator is placed above the table scan so that reordering occurs before any LLM call. This optimization is similar to a top- k operation in information retrieval, but allows the query engine to potentially process *all* tuples in the relation without enforcing a set limit. Again, computing relevance scores and sorting are inexpensive compared to LLM evaluation.

Estimation with Confidence Sequences. In addition to deterministic early stopping, EVERGREEN uses statistical estimation to resolve certain comparisons. Our approach is similar to early work on online aggregation [33], which uses classical confidence intervals (CIs). However, using CIs for online aggregation suffers from the “peeking problem” [38], where continuously monitoring the CI inflates the error rate. To avoid this issue, we instead leverage recent work on confidence sequences (CSs) [79].

More formally, let $(Z_s)_{s \geq 1}$ be a sequence of Bernoulli random variables, where $Z_s = 1$ if the s -th tuple satisfies the aggregation function’s boolean expression expr and $Z_s = 0$ otherwise. Further let $\mu \in [0, 1]$ denote the true proportion of tuples satisfying expr in an accumulator’s input stream. The accumulator estimates μ with an *anytime-valid* $(1 - \alpha)$ -confidence sequence: a sequence of confidence intervals $([L_s, U_s])_{s \geq 1}$ that satisfies

$$\Pr(\forall s \geq 1 : \mu \in [L_s, U_s]) \geq 1 - \alpha,$$

where $[L_s, U_s] \subseteq [0, 1]$ for all $s \geq 1$ and $\alpha \in (0, 1)$ is a configurable significance level. Equivalently, $1 - \alpha$ is the confidence level. Because the preceding guarantee holds simultaneously for *all* sample sizes s , the accumulator can check $[L_s, U_s]$ after each observed tuple and stop as soon as the verdict is clear, without repeated checks inflating the error rate above α . A classical CI, by contrast, is valid only at a single, predetermined sample size, so continuously monitoring it is unsound. CSs enable EVERGREEN to stop confidently as early

as possible without committing to a sample size that may be too small to be conclusive or too large to be cost-effective.

In general, a comparison resolves once the entire CI lies on one side of the threshold or within a tolerance interval (Table 1). For `proportion()`, EVERGREEN compares the proportion CI $[L_s, U_s]$ directly against ρ ; for `count_if()`, it first scales the proportion CI by the total number n of input tuples to obtain the count CI $[nL_s, nU_s]$ and compares against k . Equality comparisons additionally use a configurable relative error tolerance $\varepsilon \in (0, 1)$. For `bool_and()`, the accumulator checks whether $\mu = 1$ using the equality condition with $\rho = 1$, which reduces to $[L_s, U_s] \subseteq [1 - \varepsilon, 1]$, and therefore $L_s \geq 1 - \varepsilon$ since $U_s \leq 1$ by construction. For `bool_or()`, since a witness already triggers deterministic stopping, the scaled count CI serves mainly to detect that no satisfying tuple exists, i.e., $nU_s < 1$. When the total number n of input tuples is known, we use sampling without replacement to yield tighter CIs. EVERGREEN computes CSs via the betting approach of Waudby-Smith and Ramdas [79].

Since CSs require that observations arrive in exchangeable (i.e., uniformly random) order, the optimizer inserts a shuffle operator when estimation is used. For ungrouped aggregates, the shuffle is placed above the table scan, randomizing all tuples. For grouped aggregates, the shuffle is inserted above the sort operator used for grouping and hierarchically randomizes at each level: groups are shuffled as contiguous blocks, and tuples within the finest groups can also be shuffled if estimation is used at that aggregate level, preserving group contiguity for the streaming aggregate.

When multiple estimation decisions occur in the same query, EVERGREEN controls the *family-wise error rate*: the probability of *any* CI failing to contain its true proportion must not exceed α . The error budget is divided in two stages. First, for a query with o aggregate operators that use estimation, the budget is split equally, allocating a significance level of α/o to each operator. Second, within each operator, the budget is divided across its a estimating accumulators and the m groups in the input stream. When the number of groups m is known in advance from a sort operator, EVERGREEN applies a Bonferroni correction, assigning each accumulator a per-group significance level of $\alpha/(o \cdot a \cdot m)$. When m is not known in advance, EVERGREEN allocates a geometrically decreasing share to each successive group: each accumulator’s i -th group receives a significance level of $\alpha/(o \cdot a \cdot 2^i)$. Because $\sum_{i=1}^{\infty} 1/2^i = 1$, the total error budget never exceeds α regardless of the number of groups.

Interaction Between Sorting and Estimation. Relevance sorting and estimation impose conflicting requirements on tuple order. Relevance sorting arranges tuples to accelerate deterministic early stopping, while estimation requires exchangeable order for the CS guarantee. The optimizer thus applies at most one strategy per aggregate operator. The choice is guided by a heuristic that assumes claims are likely *true*, which we expect to hold increasingly as LLMs improve. Under this assumption, the deciding question is whether a few satisfying tuples suffice to confirm the claim: if so, relevance sorting surfaces them early; otherwise, the optimizer uses estimation. Relevance sorting therefore applies to $\exists$, $\exists_{\geq k}$, and $\exists_{>k}$ claims with small k ; estimation is unhelpful here, since tight CIs require observing more tuples. The optimizer uses estimation for all other cases, where surfacing satisfying tuples early does not accelerate confirmation. A true universal claim has $\mu = 1$, which estimation confirms from a sample without processing every




Figure 4: Optimized query plan for the verification query in Figure 2.

tuple. For $\exists_{\geq k}$ and $\exists_{> k}$ claims with large k , sorting does not pay off because too many satisfying tuples are needed. $\exists_{\leq k}$, $\exists_{< k}$, and $\exists_{=k}$ claims require bounding the total number of satisfying tuples, which entails processing nearly all tuples regardless of order. A true $\exists_{\neq k}$ claim’s count may lie above or below k ; relevance sorting accelerates only the former by surfacing $k + 1$ satisfying tuples, whereas estimation can stop early in either case. For proportional claims, the optimizer cannot determine whether the equivalent count threshold $n\rho$ is small, since n is unknown at optimization time. In queries with nested aggregates, the two strategies can coexist at different levels: the inner aggregate uses relevance sorting within each group, while the outer aggregate uses estimation with group order shuffled. We note that this bias toward true claims can backfire. For example, if an existential claim is false, relevance sorting calls an LLM on the entire relation, while estimation would have been more efficient. Adaptively selecting between relevance sorting and estimation at runtime is a promising future direction.

End-to-End Example. We illustrate how EVERGREEN optimizes the query in Figure 2, yielding the plan in Figure 4. Before any LLM call, the plan reorders tuples: a relevance sort surfaces reviews likely to complain about poor service quality, and a sort on `business_id` makes each location’s reviews contiguous so `count_if()`’s streaming aggregate can process one group at a time. A shuffle then randomizes group order—preserving within-group relevance—so the outer `bool_and()`’s CS observes groups in exchangeable order. The semantic `map()` labels only the reviews pulled through. The two aggregates apply different stopping conditions from Table 1: the inner `count_if()` can only stop deterministically, while the outer `bool_and()` can stop both deterministically and probabilistically.

4.2 General-Purpose Optimizations

Operator Fusion. A verification query can contain multiple semantic operators, e.g., a semantic filter followed by a semantic map, each requiring a separate LLM call per tuple. EVERGREEN reduces this cost by fusing consecutive operators, combining their prompts into a single LLM call. If any filter evaluates to false, the fused operator discards the tuple; otherwise, it emits the map outputs.

When the fused operator contains a semantic filter, the cardinality scan operator—which provides an aggregate with the number n of input tuples—is no longer used. This is because neither placement of the cardinality scan works: placing it *below* the fused operator yields a pre-filter count that *overstates* the number of tuples the aggregate will observe, while placing it *above* executes the fused operator’s LLM on *every* tuple before the aggregate begins (since it is a pipeline breaker), defeating the purpose of early stopping. Without n , three capabilities described in Section 4.1 become unavailable: (1) deterministic bounds based on the maximum achievable count or proportion, (2) scaling the proportion CI to a count CI

for `count_if()`, and (3) tighter CIs via sampling without replacement. In such cases, the optimizer falls back to other optimizations that do not require knowing n . For instance, the inability to scale the CI prevents estimation for $\exists_{\geq k}$ and $\exists_{> k}$ cardinal claims, so the optimizer applies relevance sorting regardless of the threshold k . Even though the stopping conditions that require n are unavailable, operator fusion still reduces the total number of LLM calls.

Fusing multiple prompts into a single LLM call may reduce output quality. Yet, based on our observations, verification queries usually require only a few narrowly scoped semantic operators. Given the current capabilities of LLMs, fusion thus reduces the number of LLM calls with minimal quality degradation.

Similarity Filtering. Semantic filters invoke an LLM on every input tuple, even though many may not satisfy the predicate. EVERGREEN avoids these calls by inserting a similarity pre-filter below each semantic filter. At optimization time, the optimizer reuses the same LLM-driven approach as in relevance sorting, generating a search query from the filter’s prompt and embedding it. At execution time, the pre-filter discards a tuple if the maximum cosine similarity between the query embedding and the tuple’s sentence-level embeddings falls below a configurable threshold. Sentence-level embeddings avoid the signal dilution of a single document embedding, retaining a tuple when even one of its sentences is relevant. A conservatively low threshold keeps recall high while still eliminating tuples clearly unrelated to the predicate.

Prompt Caching. LLM responses are cached on disk, keyed by model name and prompt string. When the same prompt is sent to a model, the cached response is reused to avoid redundant LLM evaluation. This benefits workloads where multiple claims share semantic operators over the same underlying data.

End-to-End Example. The query in Figure 3 exercises the general-purpose optimizations. A similarity pre-filter discards reviews unrelated to service before LLM evaluation, and the semantic filter() and map() are fused into a single LLM call per review. Unlike the query in Figure 2, ranking requires the full `proportion()` of each location, so the aggregate does not stop early; the savings instead come from issuing fewer LLM calls per tuple. Finally, when related ordinal claims are verified (e.g., which location ranks first vs. second in service), prompt caching reuses LLM responses across claims.

5 CITATIONS AS PROVENANCE

Returning a boolean verdict is often insufficient for users; each verdict should be paired with a collection of citations so that users can check its underlying evidence. By treating claims as queries, EVERGREEN formalizes citations as *data provenance*, giving users a precise explanation of the verdict based on facts in the relation. Specifically, we ground our work in *semiring provenance for first-order logic* [31]. We first provide a short overview of this formalization and then describe its application and extension to our setting.

5.1 Semiring Provenance for First-Order Logic

Seminal work on provenance semirings for *positive relational algebra* [30] annotates each output tuple with a *polynomial* over a set of indeterminates X in a *commutative semiring*. Each *indeterminate* $p_t \in X$, also called a *provenance token*, uniquely annotates an input tuple $t \in \mathcal{R}$. The semiring’s *addition* (+) and *multiplication* (·) represent the *alternative* (e.g., union) and *joint* (e.g., join) use of input tuples, respectively. Consequently, each *monomial* in the polynomial is an *alternative explanation* of which input tuples jointly contribute to the output tuple. Grädel and Tannen [31] extend this framework to full *first-order logic with negation* by pairing each indeterminate $p_t \in X$ (a *positive token*) with a *dual indeterminate* $\bar{p}_t \in \bar{X}$ (a *negative token*). In this semiring of *dual-indeterminate polynomials*, addition corresponds to disjunction and existential quantification, while multiplication corresponds to conjunction and universal quantification. In our setting, a positive token p_t records the *positive fact* that tuple t satisfies φ (i.e., $\varphi(t) = \top$), while a negative token $\bar{p}_t$ records the *negative fact* that t does *not* satisfy φ (i.e., $\varphi(t) = \perp$). For example, given $\varphi :=$ “enjoys the restaurant’s chicken salad”, p_t records that tuple t enjoys the restaurant’s chicken salad, while $\bar{p}_{t'}$ records that another tuple t' does *not* enjoy it.

More formally, an *interpretation* π maps each literal to a polynomial and extends *inductively* to all first-order formulas with the following rules (adapted from Definition 2 of [31]): (i) $\pi[\psi \vee \chi] := \pi[\psi] + \pi[\chi]$, (ii) $\pi[\psi \wedge \chi] := \pi[\psi] \cdot \pi[\chi]$, (iii) $\pi[\exists t \in \mathcal{R} : \psi(t)] := \sum_{t \in \mathcal{R}} \pi[\psi(t)]$, (iv) $\pi[\forall t \in \mathcal{R} : \psi(t)] := \prod_{t \in \mathcal{R}} \pi[\psi(t)]$, and (v) $\pi[\neg\psi] := \pi[\text{nnf}(\neg\psi)]$, where ψ and χ are formulas and $\text{nnf}(\cdot)$ denotes the transformation to negation normal form. In our setting, we treat each $\varphi(t)$ and its negation $\neg\varphi(t)$ as literals regardless of φ ’s internal structure, so $\pi[\varphi(t)] = \pi(\varphi(t))$ and $\pi[\neg\varphi(t)] = \pi(\neg\varphi(t))$, which form the base cases for the inductive interpretation. Since we want to track both positive and negative facts, we set $\pi(\varphi(t)) = p_t$ when $\varphi(t) = \top$ and $\pi(\varphi(t)) = 0$ otherwise; dually, $\pi(\neg\varphi(t)) = \bar{p}_t$ when $\varphi(t) = \perp$ and $\pi(\neg\varphi(t)) = 0$ otherwise. That is, 0 is interpreted as a *false assertion*.

Grädel and Tannen’s main contribution [31] is support for *negation*. While the interpretation π produces a polynomial explaining a *true* sentence, it yields 0 for a *false* one, giving no explanation. In this case, π can instead be applied to the sentence’s negation, which is true, revealing why the original is false. This framework thus provides explanations for both cases, i.e., *why-provenance* and *why-not provenance*, so EVERGREEN can justify each verdict, whether true or false. We give examples of this approach in the next subsection.

5.2 Provenance Semantics for Claims

EVERGREEN derives its provenance from the prior interpretation rules. However, rather than computing the full provenance polynomial (i.e., all alternative explanations), EVERGREEN returns a single monomial (i.e., one explanation) either because the polynomial is itself a monomial or because of optimizations like early stopping. By the semantics of the claim types, every complete monomial EVERGREEN observes is *minimal* (i.e., irredundant): no proper subset is also a monomial of the full polynomial. This is immediate for true existential claims (a single token per monomial) and true universal claims (a single monomial over all tuples), and follows for the other claim types from their polynomials derived later in this section.

Table 2: Explanation associated with each claim’s verdict. $W \subseteq \mathcal{R}$ is a set of tuples satisfying φ . $\exists_{>k}$ follows from $\exists_{\geq k+1}$. $\exists_{<k}$, $\exists_{\leq k}$, and $\exists_{\neq k}$ swap true/false explanations of $\exists_{\geq k}$, $\exists_{>k}$, and $\exists_{=k}$, respectively. Proportional claims reduce to the cardinal rows; ordinal and nested claims compose these rows.

	Explanation for True ($\top$)	Explanation for False ($\perp$)
$\exists$	p_t (first witness)	$\prod_{t \in \mathcal{R}} \bar{p}_t$
$\forall$	$\prod_{t \in \mathcal{R}} p_t$	$\bar{p}_t$ (first counterexample)
$\exists_{\geq k}$	$\prod_{t \in W} p_t, W = k$	$\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}, W < k$
$\exists_{\leq k}$	$\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}, W = k$	$\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}, W \neq k$

When the verdict is instead settled before a complete monomial is observed (e.g., estimation over a sample $S \subset \mathcal{R}$ or deterministic early stopping from a maximum achievable count bound), the returned tokens form a *subset* of such a minimal monomial. Table 2 summarizes the explanation EVERGREEN returns for each claim type and verdict in complete monomial form, which we now derive.

Existential and Universal Claims. For a *true existential* claim $\exists t \in \mathcal{R} : \varphi(t)$, applying rule (iii) yields $\sum_{t \in \mathcal{R}} \pi(\varphi(t))$. Each tuple t contributes $\pi(\varphi(t)) = p_t$ to the polynomial if $\varphi(t) = \top$ and $\pi(\varphi(t)) = 0$ otherwise. Since EVERGREEN stops early and any monomial suffices as a *witness*, the returned provenance is the corresponding positive token p_t for the first tuple t where $\varphi(t) = \top$. For example, given the existential claim “Some reviewers enjoy the restaurant’s chicken salad”, if t is the first tuple that satisfies $\varphi :=$ “enjoys the restaurant’s chicken salad”, then the provenance is p_t .

For a *false existential* claim, we first apply rule (v) to interpret its *negation* $\forall t \in \mathcal{R} : \neg\varphi(t)$. Applying rule (iv) then gives $\prod_{t \in \mathcal{R}} \pi(\neg\varphi(t)) = \prod_{t \in \mathcal{R}} \bar{p}_t$. In this case, EVERGREEN returns this single monomial containing all negative tokens, which indicates that no tuple in $\mathcal{R}$ satisfies φ (e.g., no tuple enjoys the chicken salad). When estimation is enabled and $nU_s < 1$, EVERGREEN returns only $\prod_{t \in S} \bar{p}_t$ for the processed sample $S \subset \mathcal{R}$.

The universal cases are the dual of the existential cases. For a *true universal* claim $\forall t \in \mathcal{R} : \varphi(t)$, applying rule (iv) yields $\prod_{t \in \mathcal{R}} \pi(\varphi(t)) = \prod_{t \in \mathcal{R}} p_t$. Hence, EVERGREEN returns this monomial of all positive tokens, which indicates that each tuple in $\mathcal{R}$ satisfies φ . When estimation is enabled and $L_s \geq 1 - \varepsilon$, EVERGREEN returns only $\prod_{t \in S} p_t$ for the processed sample.

For a *false universal* claim, we apply rule (v) to interpret its *negation* $\exists t \in \mathcal{R} : \neg\varphi(t)$. Applying rule (iii) then gives $\sum_{t \in \mathcal{R}} \pi(\neg\varphi(t))$. Each tuple t contributes $\pi(\neg\varphi(t)) = \bar{p}_t$ if $\varphi(t) = \perp$ and $\pi(\neg\varphi(t)) = 0$ otherwise. Since EVERGREEN stops early and any monomial suffices as a *counterexample*, the returned provenance is the corresponding negative token $\bar{p}_t$ for the first tuple t where $\varphi(t) = \perp$.

Cardinal and Proportional Claims. Grädel and Tannen [31] do not explicitly provide provenance expressions for cardinal quantifiers, so we derive them based on their interpretation rules.

First, $\exists_{\geq k} t \in \mathcal{R} : \psi(t)$ is equivalent to k nested existential quantifiers over strictly ordered tuples, each satisfying ψ . Interpreting in the semiring replaces each existential quantifier with a sum and each conjunction with a product, yielding

$$\pi[\exists_{\geq k} t \in \mathcal{R} : \psi(t)] = \sum_{\substack{W \subseteq \mathcal{R} \\ |W|=k}} \prod_{t \in W} \pi[\psi(t)].$$

Setting $\psi = \varphi$, if $\varphi(t) = \top$ for all $t \in W$, then $\prod_{t \in W} \pi(\varphi(t)) = \prod_{t \in W} p_t \neq 0$; otherwise, the product is 0. The polynomial is thus the sum of monomials, one for each size- k combination of positive tokens. The polynomial for $\exists_{>k}$ simply replaces k with $k + 1$.

In contrast, $\exists_{=k} t \in \mathcal{R} : \psi(t)$ is equivalent to the same k nested existential quantifiers, but additionally requires that every remaining tuple does *not* satisfy ψ . Interpreting in the semiring augments the $\exists_{\geq k}$ polynomial with a product over the remaining tuples $t' \in \mathcal{R} \setminus W$, each contributing $\pi[\neg\psi(t')]$, yielding

$$\pi[\exists_{=k} t \in \mathcal{R} : \psi(t)] = \sum_{\substack{W \subseteq \mathcal{R} \\ |W|=k}} \left(\prod_{t \in W} \pi[\psi(t)] \prod_{t' \in \mathcal{R} \setminus W} \pi[\neg\psi(t')] \right).$$

Setting $\psi = \varphi$, if $\varphi(t) = \top$ for all $t \in W$ and $\varphi(t') = \perp$ for all $t' \in \mathcal{R} \setminus W$, then we have that $\prod_{t \in W} \pi(\varphi(t)) \prod_{t' \in \mathcal{R} \setminus W} \pi(\neg\varphi(t')) = \prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'} \neq 0$; otherwise, the product is 0. Since each tuple either satisfies φ or does not, at most one subset W yields a nonzero product, precisely when $|W| = k$, i.e., the sentence holds. The polynomial is then a single monomial $\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}$: a positive token for each of the k tuples satisfying φ and a negative token for the remaining $|\mathcal{R}| - k$. The provenance polynomials for $\exists_{<k}$, $\exists_{\leq k}$, and $\exists_{\neq k}$ are the same as $\exists_{=k}$ but sum over subsets W of size $< k$, $\leq k$, and $\neq k$, respectively. As with $\exists_{=k}$, requiring every tuple outside W not to satisfy φ forces W to be exactly the set of satisfying tuples; every other subset contributes 0. Hence, the polynomials for the other cardinal quantifiers also reduce to at most the single monomial $\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}$, where W is the satisfying set. The size constraint on W ($< k$, $\leq k$, or $\neq k$) only determines whether this monomial is included in the sum, i.e., whether the polynomial is nonzero and thus the sentence holds.

Given these polynomials, we now describe the provenance EVERGREEN returns for each cardinal claim type. For a *true cardinal* claim $\exists_{\geq k} t \in \mathcal{R} : \varphi(t)$, at least one monomial $\prod_{t \in W} p_t$ of $|W| = k$ positive tokens exists. Since early stopping halts after finding k satisfying tuples, EVERGREEN returns only these first k positive tokens. For example, if the claim “At least a handful of vegetarian customers enjoyed the restaurant’s burgers” is true, EVERGREEN returns the first k positive tokens with tuples satisfying $\varphi :=$ “vegetarian customer who enjoyed the restaurant’s burgers”. Similar to the prior claim types, when estimation is enabled and $nL_s \geq k$, EVERGREEN returns only the positive tokens in the processed sample $S \subset \mathcal{R}$.

For a *false cardinal* claim $\exists_{\geq k} t \in \mathcal{R} : \varphi(t)$, its *negation* $\exists_{<k} t \in \mathcal{R} : \varphi(t)$ is true, so its provenance is $\prod_{t \in W} p_t \prod_{t' \in \mathcal{R} \setminus W} \bar{p}_{t'}$ where $|W| < k$, which demonstrates that fewer than k tuples satisfy φ . When only a subset of tuples is processed—either because the count provably cannot reach k given the remaining tuples, or because $nU_s < k$ —the tokens cover only the processed tuples.

For $\exists_{=k}$ and $\exists_{\neq k}$, when the count equals k , EVERGREEN processes all tuples and outputs all tokens to demonstrate the exact count; when the count differs from k , it may stop early and output the processed tokens, demonstrating that the count lies above or below k . Since proportional claims reduce to cardinal claims, EVERGREEN scales the proportion threshold ρ by the number of processed tuples and then applies the provenance semantics of cardinal claims.

Ordinal Claims. An ordinal claim asserts that a target group G_i achieves rank r among m groups, where the rank is based on a per-group aggregate f . This is determined by a conjunction of $m - 1$ pairwise cardinal comparisons between G_i and every other group G_j ($i \neq j$); in the semiring, this conjunction becomes a product of per-comparison polynomials. Each comparison produces two monomials (one per group) following the provenance semantics for cardinal or proportional claims depending on the aggregate f . When $f(G_i, \varphi) > f(G_j, \varphi)$, G_i ’s tokens follow the rules for a true $\exists_{>f(G_j, \varphi)}$ claim, while G_j ’s tokens follow the rules for a true $\exists_{<f(G_i, \varphi)}$ claim. When $f(G_i, \varphi) < f(G_j, \varphi)$, the roles are reversed. When $f(G_i, \varphi) = f(G_j, \varphi)$, both groups’ tokens follow the rules for a true $\exists_{=f(G_i, \varphi)}$ claim and contribute all tokens.

Nested Claims. A nested claim composes an outer quantifier over groups with an inner claim per group. Such claims are already supported by the *inductive* interpretation rules. EVERGREEN first evaluates each group’s inner claim, producing per-group provenance tokens following the semantics of the inner claim type. The outer quantifier then combines these per-group tokens following its provenance semantics. For example, if the claim $\forall G \in \mathcal{G} : (\exists_{\geq 2} t \in G : \varphi(t))$ is true, the provenance is the product of two positive tokens per group. Otherwise, the provenance is the first failing group’s tokens, showing that fewer than two of its tuples satisfy φ .

6 EXPERIMENTAL EVALUATION

Our experimental evaluation aims to show that EVERGREEN achieves higher verification quality at lower cost than existing approaches. Beyond this end-to-end comparison, we analyze the quality of EVERGREEN’s provenance and semantic operators, and ablate the effect of each optimization on quality and cost.

Benchmark Datasets and Claims. To the best of our knowledge, no benchmark targets claim verification over semantic aggregates, so we curate one to evaluate EVERGREEN. From our analysis of production semantic aggregation queries on Snowflake AISQL [48], we observe that queries are often executed over relations of customer reviews and support tickets. As such, we first select three subsets of the restaurant reviews in the Yelp Open Dataset [84], which we identify as *johns* (1,609 tuples; 230k tokens), *mcdonalds* (1,813 tuples; 244k tokens), and *village* (1,603 tuples; 291k tokens). *mcdonalds* contains reviews from 62 McDonald’s locations in Missouri, while *johns* and *village* are reviews of separate restaurants. We then include three subsets of the Customer Support on Twitter dataset [10], which we call *play* (1,197 tuples; 394k tokens), *airlines* (3,167 tuples; 1,213k tokens), and *uber* (2,687 tuples; 1,005k tokens). *airlines* contains dialogs between customers and support agents across six airlines, while *play* and *uber* contain dialogs for PlayStation and Uber, respectively. We follow prior work [25] to preprocess the dialogs, keeping only those that contain 6–20 utterances between a single customer and agent.

We pose three types of aggregation queries over the six datasets via AISQL’s `AI_AGG()` function [48] with Llama 3.3 70B [29]. We issue separate summarization queries to *johns*, *village*, *play*, and *uber* (e.g., “Summarize what customers are saying about John’s Roast Pork”). For *mcdonalds* and *airlines*, we issue comparative queries (e.g., “Compare the reviews for the different McDonald’s locations; highlight the commonalities and differences between

Table 3: Representative sample of benchmark claims.

ID	Dataset	Semantic Agg. Type	Claim Type	Claim	Grounded
C2	johns	Summarize	$\exists$	Some reviewers enjoy the restaurant's chicken salad.	T
C3			$\exists_{\geq 0.1}$	Common criticisms of John's Roast Pork include cash-only policy.	$\perp$
C4			$\forall$	The reviews do not mention any vegan options.	T
C8	mcdonalds	Compare	$\forall \exists_{\geq 2}$	All McDonald's locations have multiple complaints about poor service quality.	T
C9		Rank	$\text{rank}() = 1$	The top-ranked McDonald's location in terms of service has the Business ID [A].	T
C10			$\text{rank}() = 2$	The McDonald's location ranked #2 in terms of service has the Business ID [B].	T
C16	village	Summarize	$\forall$	There are over 200 varieties of whiskey available at Village Whiskey.	$\perp$
C19	play	Summarize	$\exists_{=0.1}$	10% of customers are experiencing problems with payments or billing.	$\perp$
C21	airlines	Compare	$\exists_4 \exists_{>0.15}$	There are 4 airline companies where over 15% of customer complaints are regarding flight delays.	T
C26		Rank	$\text{rank}() = 2$	SouthwestAir has the second most customer support dialogs with complaints about flight booking issues.	T
C29		uber	Summarize	$\exists_{=776}$	776 customers complained about poor driver behavior.
C30	$\exists_{<725}$			Less than 725 customers complained about poor driver behavior.	$\perp$
C31	$\exists_{<0.1}$			Less than 10% of customers reported receiving inconsistent or inaccurate information from support agents.	$\perp$
C32	$\exists_{<0.25}$			Less than 25% of customers reported receiving inconsistent or inaccurate information from support agents.	T

the restaurant locations”) and ranking queries (e.g., “Analyze the customer support dialogs for the different airline companies and describe the worst 3 airline companies regarding flight booking issues”). We use the approach described in Section 2 to decompose each aggregate response into a list of claims.

To curate a diverse and balanced benchmark, we start from these automatically decomposed claims. We use some directly, adapt others (e.g., recasting a reported proportion as a count, or negating a claim to obtain a false variant), and add manually written claims. In total, we obtain 32 claims that span all claim types described in Section 2, with an even split between grounded and ungrounded claims. Table 3 shows a representative sample. In general, complete human annotation of claim-level ground truth is infeasible in our setting because verifying quantified claims (e.g., “all reviews are positive”) may require reviewing every tuple in the relation—over 1,000 per dataset. Instead, the ground truth validity of each claim is determined with a reference implementation of EVERGREEN that executes manually written verification queries. Each invocation of a semantic operator takes the majority vote from an ensemble of three strong LLMs spanning distinct families (Claude Opus 4.8 [8], GPT-5.5 [57], and Gemini 3.5 Flash [28]). Moreover, the reference implementation disables all optimizations, since some can degrade quality and early stopping prevents capturing tuple-level labels for provenance and semantic operator quality evaluations. To evaluate agreement of the reference labels against an independent human reference, we first draw a stratified random sample of per-operator decisions (i.e., the ensemble’s tuple-level filter and map outputs) covering all 26 distinct semantic operators in our benchmark. We then sample four decisions per operator, balanced across label values, for a total of 104 decisions. The first author, blind to the ensemble labels, annotates each sampled decision, and we measure agreement between the human and ensemble using exact match and Cohen’s κ . The human and ensemble agree on 98.1% of decisions (95% Wilson CI [0.93, 0.99]; $\kappa = 0.96$). As an additional check, for claims whose verdict can be determined via a handful of tuples (e.g., $\exists$ or $\exists_{\geq k}$ with small k), we manually confirmed the ensemble’s verdict: these checks agreed with the ensemble in every case.

Implementations. We use Snowflake Cortex AI [70] for LLM inference and embedding generation. All embeddings are created by

Snowflake arctic-embed-l-v2.0 [53]. We compare EVERGREEN with four implementations: `base_rm`, `rag_agent`, `r1m`, and `evg_unopt`.

base_rm prompts an LLM to reason about whether a claim is supported by the given input relation and represents existing LLM-as-a-judge baselines [54, 72, 91]. We instantiate `base_rm` with three LLMs: Claude Opus 4.6 [7], Sonnet 4.6 [9], and Haiku 4.5 [5]. It serializes each tuple as JSON and greedily fills the prompt up to Claude’s 200k-token context window. When claims involve vague quantifiers (e.g., “common”), we provide hints specifying exact thresholds (e.g., “at least 10%”) for a fairer comparison with the reference ground truth. These hints are also used by the other implementations. We set temperature to 0 and allow up to 10 retries, since `base_rm` occasionally exhausts its 8,192-token output limit.

rag_agent is an LLM agent that uses a retrieval tool and represents an implementation of iterative retrieval-based claim verification approaches [41, 83]. The interface to the retrieval tool is similar to relevance sorting but restricted to top- k queries. In addition to a semantic search query, inclusion keywords, and exclusion keywords, the agent must specify a limit on the most relevant tuples to retrieve from the database. The agent can also include structured filters in the retrieval query (e.g., `business_id = [A]`), which are executed as pre-filters before retrieving the requested number of tuples. After tuples are retrieved, the LLM can choose the number of tuples to view initially from the set $\{10, 20, \dots, 50\}$. At each iteration of the agentic loop, the LLM either (1) issues a retrieval query and views the initial results, (2) views the next page of previously retrieved tuples, or (3) outputs whether the claim is grounded and terminates. We instantiate `rag_agent` with the same three Claude models as `base_rm` and set temperature to 1.

r1m is a Recursive Language Model (RLM) [87] that can iteratively write and execute arbitrary Python code in a REPL, where the generated code can also invoke sub-LLM calls. We additionally provide `r1m` with the same retrieval tool given to `rag_agent`. As such, this implementation is a generalization of prior imperative, program-based verification approaches [34, 58]. We instantiate `r1m` using the recent implementation in DSPy [42] and keep the root LLM fixed to Opus while varying the sub-LLM to the same three Claude models used for `base_rm` and `rag_agent`. The root LLM uses a temperature of 1 while the sub-LLM uses a temperature of 0.

evg_unopt instantiates EVERGREEN with no optimizations enabled and represents a naive semantic query processing engine like BINDER [18]. It uses each of the three Claude models from earlier and evaluates semantic operators on every tuple. All EVERGREEN implementations use a temperature of 0. Moreover, implementations process tuples in batches of 32, trading up to a batch’s worth of unnecessary LLM calls for significantly lower latency.

evg_opt instantiates EVERGREEN with all optimizations enabled. We evaluate five configurations, each using a different LLM for its semantic operators. We include the three Claude models used in the other implementations and two additional Qwen models (Qwen3-VL-235B-A22B [11] and Qwen3-Next-80B-A3B [62]) to evaluate how quality and cost change with even weaker models. Opus is always used for the optimizer’s LLM to generate high-quality search queries. To maintain high recall, we set the threshold for similarity filters to 0.15. For estimation with CSs, we set $\alpha = \epsilon = 0.05$.

For EVERGREEN implementations, claims are compiled to queries by providing an LLM with the claim, the aggregation prompt, the dataset schema, hints for vague quantifiers, and EVERGREEN’s API documentation (with example queries from a separate movie review domain). Averaged over the 32 claims and three trials each, it costs $\$0.051 \pm 0.001$ and takes $9.8 \pm 1.2s$ (mean $\pm$ SD) to compile a claim using Opus 4.6, with negligible variation across claim types. We reviewed every generated query and found each equivalent to its manually written reference query. We attribute this reliable compilation to EVERGREEN’s declarative, constrained query interface.

Verification Results. Figure 1 highlights the cost–quality trade-offs for each implementation, and Table 4 reports their quality, cost, and latency in more detail. Precision, recall, and F1 score treat ungrounded claims as the positive class, measuring an implementation’s ability to detect hallucinations. Since our benchmark has an even split between grounded and ungrounded claims, we also report accuracy. Cost is computed from input and output token counts at published per-million-token USD rates for each LLM. We use Anthropic’s input/output rates [6] of $\$5/\25 for Opus 4.6, $\$3/\15 for Sonnet 4.6, and $\$1/\5 for Haiku 4.5, and Alibaba Cloud’s rates [2] of $\$0.287/\1.147 for Qwen3-VL and $\$0.144/\0.574 for Qwen3-Next. Since latency is subject to Cortex’s server-side configuration and variability, we treat it as a secondary metric.

evg_opt configurations form the entire Pareto frontier: all other implementations are dominated by some evg_opt point with higher or equal F1 at lower cost. Compared to evg_unopt with Opus, evg_opt preserves verification quality at an F1 of 0.94 while reducing cost by $3.1\times$ ($\$5.05$ vs. $\$15.88$), demonstrating the effectiveness of EVERGREEN’s optimizations. rlm with Sonnet performs the best among the non-EVERGREEN implementations with an F1 of 0.83; however, evg_opt with Qwen3-VL dominates at 0.87 with $7.0\times$ lower cost ($\$0.183$ vs. $\$1.29$). These results highlight the advantages of EVERGREEN’s declarative, neuro-symbolic design. By enabling the LLM compiler to simply specify the high-level verification query, EVERGREEN abstracts away optimization details and restricts LLMs to fine-grained semantic tasks. This yields higher verification quality at lower cost with weaker models than approaches that rely on strong LLMs for both semantic and symbolic processing or program-based approaches that push optimization onto the LLM.

Figure 5 reveals each implementation’s failure modes. base_rm fails when grounding witnesses fall beyond its context window (e.g.,

a review praising the restaurant’s chicken salad for existential claim C2, the target entities for ordinal claims C10 and C26, and all but one airline for nested claim C21). base_rm also struggles with precise quantities, defaulting to ungrounded on cardinal claim C29 and proportional claim C32. rag_agent avoids the first problem since retrieval surfaces the relevant witnesses, but because it estimates quantities in natural language from retrieved samples, it still fails on claims requiring precise aggregates (e.g., C10, C19, C26, and C29). While rlm improves over rag_agent, it sometimes processes only a sample of the relation with its sub-LLM, so it faces similar sampling imprecision (e.g., cardinal claim C30 and proportional claim C31). Even when rlm processes every tuple, noisy per-tuple labels and small per-group denominators flip the ranking for ordinal claim C9. In contrast, EVERGREEN’s only source of error is semantic operator mislabeling, as its symbolic aggregation is otherwise exact or statistically sound. Errors occur more frequently with weaker models; for example, on C30, only Opus returns the correct verdict. However, claims involving close rankings (C10) or exact counts (C29) remain challenging even for the strongest model.

Provenance and Semantic Operator Quality. To understand how EVERGREEN’s provenance and semantic operators degrade with weaker models, we execute the reference queries with evg_unopt and evg_opt across LLMs and measure provenance precision and per-operator quality (Table 5). Semantic filter and map quality are micro-averaged over distinct semantic operators, while provenance precision is pooled across all queries. Provenance precision is the proportion of returned tokens that are correct according to the reference implementation. We report precision rather than recall because EVERGREEN intentionally returns non-exhaustive provenance tokens; we therefore require provenance to be correct, not complete. We measure filter quality by precision, recall, and F1, treating tuples selected by the reference’s filter as the positive class. Map accuracy is the proportion of map outputs that match the reference. evg_opt with Opus closely tracks evg_unopt: it matches provenance precision at 0.94 and has comparable filter precision (0.99 vs. 0.98). Similarity filtering modestly reduces filter recall to 0.90 from 0.97, lowering filter F1 to 0.94 from 0.98, while map accuracy dips slightly to 0.94 from 0.96. Quality then degrades gracefully as the model weakens: filter F1 and map accuracy remain at least 0.90 and 0.88, respectively, across all models, while provenance precision is the most sensitive, dropping to 0.82 from 0.94.

Ablation Study. We evaluate the contribution of each optimization by performing a leave-one-out ablation study with Haiku, disabling one optimization while keeping the others enabled. Figure 6a reports the resulting cost multipliers relative to the fully optimized configuration; we summarize each optimization by the geometric mean of these ratios over its applicable claims. Note that disabling early stopping effectively disables relevance sorting and estimation.

Early stopping leads to the largest cost savings of $9.5\times$, illustrating the benefits of verification-aware optimizations. Estimation and operator fusion each give comparable savings of $2.9\times$ to $3.0\times$. Relevance sorting provides more moderate savings of $1.8\times$. Similarity filtering produces almost no cost savings in our evaluation because the similarity threshold is set low to maintain high recall, filtering out few tuples before the semantic filter is executed. Prompt caching is omitted from Figure 6a: four of the eight applicable claims are served entirely from cache at no cost when enabled, so their cost

Table 4: Verification quality, average cost, and average latency across claims. We report the mean [min, max] over three trials.

Implementation	LLM	Precision	Recall	F1 Score	Accuracy	Cost (\$)	Latency (s)
base_rm	Claude Opus 4.6	0.71 [0.70, 0.73]	0.94 [0.88, 1.00]	0.81 [0.78, 0.84]	0.78 [0.75, 0.81]	0.945 [0.942, 0.947]	47 [39, 58]
	Claude Sonnet 4.6	0.69 [0.67, 0.71]	0.90 [0.88, 0.94]	0.78 [0.76, 0.81]	0.75 [0.72, 0.78]	0.570 [0.568, 0.573]	92 [76, 105]
	Claude Haiku 4.5	0.61 [0.60, 0.62]	0.96 [0.94, 1.00]	0.74 [0.73, 0.76]	0.67 [0.66, 0.69]	0.188 [0.188, 0.188]	17 [17, 17]
rag_agent	Claude Opus 4.6	0.70 [0.63, 0.75]	0.75 [0.75, 0.75]	0.72 [0.69, 0.75]	0.71 [0.66, 0.75]	2.61 [2.46, 2.74]	220 [216, 225]
	Claude Sonnet 4.6	0.73 [0.67, 0.76]	0.79 [0.75, 0.81]	0.76 [0.71, 0.79]	0.75 [0.69, 0.78]	2.84 [2.57, 3.00]	324 [294, 355]
	Claude Haiku 4.5	0.61 [0.60, 0.62]	0.88 [0.75, 1.00]	0.72 [0.67, 0.76]	0.66 [0.62, 0.69]	0.272 [0.258, 0.287]	38 [33, 41]
rlm	Claude Opus 4.6	0.83 [0.80, 0.87]	0.79 [0.75, 0.81]	0.81 [0.77, 0.84]	0.81 [0.78, 0.84]	1.89 [1.77, 2.02]	260 [257, 261]
	Claude Sonnet 4.6	0.89 [0.86, 0.92]	0.79 [0.75, 0.88]	0.83 [0.80, 0.88]	0.84 [0.81, 0.88]	1.29 [1.21, 1.32]	199 [190, 212]
	Claude Haiku 4.5	0.80 [0.73, 0.86]	0.73 [0.69, 0.75]	0.76 [0.71, 0.80]	0.77 [0.72, 0.81]	0.649 [0.616, 0.680]	177 [163, 193]
evg_unopt	Claude Opus 4.6	0.89 [0.89, 0.89]	1.00 [1.00, 1.00]	0.94 [0.94, 0.94]	0.94 [0.94, 0.94]	15.88 [15.88, 15.88]	331 [319, 353]
	Claude Sonnet 4.6	0.83 [0.83, 0.83]	0.94 [0.94, 0.94]	0.88 [0.88, 0.88]	0.88 [0.88, 0.88]	9.54 [9.54, 9.54]	304 [290, 326]
	Claude Haiku 4.5	0.83 [0.83, 0.83]	0.94 [0.94, 0.94]	0.88 [0.88, 0.88]	0.88 [0.88, 0.88]	3.23 [3.23, 3.23]	354 [316, 400]
evg_opt	Claude Opus 4.6	0.89 [0.89, 0.89]	1.00 [1.00, 1.00]	0.94 [0.94, 0.94]	0.94 [0.94, 0.94]	5.05 [4.86, 5.36]	261 [193, 317]
	Claude Sonnet 4.6	0.88 [0.88, 0.88]	0.94 [0.94, 0.94]	0.91 [0.91, 0.91]	0.91 [0.91, 0.91]	3.25 [3.17, 3.37]	192 [169, 227]
	Claude Haiku 4.5	0.83 [0.83, 0.83]	0.94 [0.94, 0.94]	0.88 [0.88, 0.88]	0.88 [0.88, 0.88]	1.01 [1.00, 1.01]	133 [103, 179]
	Qwen3-VL	0.83 [0.82, 0.83]	0.92 [0.88, 0.94]	0.87 [0.85, 0.88]	0.86 [0.84, 0.88]	0.183 [0.174, 0.188]	228 [216, 234]
	Qwen3-Next	0.82 [0.81, 0.82]	0.83 [0.81, 0.88]	0.82 [0.81, 0.85]	0.82 [0.81, 0.84]	0.113 [0.111, 0.116]	114 [111, 115]




Figure 5: Per-claim verification correctness. Each cell shows the number of correct trials out of three.

Table 5: EVERGREEN’s provenance precision and semantic operator quality. We report the mean [min, max] over three trials.

Implementation	LLM	Provenance Precision	Filter Precision	Filter Recall	Filter F1 Score	Map Accuracy
evg_unopt	Claude Opus 4.6	0.94 [0.94, 0.94]	0.98 [0.98, 0.98]	0.97 [0.97, 0.98]	0.98 [0.98, 0.98]	0.96 [0.96, 0.96]
	Claude Sonnet 4.6	0.91 [0.91, 0.91]	0.98 [0.98, 0.98]	0.97 [0.97, 0.97]	0.97 [0.97, 0.97]	0.95 [0.95, 0.95]
	Claude Haiku 4.5	0.83 [0.83, 0.83]	0.93 [0.92, 0.93]	0.99 [0.99, 0.99]	0.96 [0.95, 0.96]	0.93 [0.93, 0.93]
evg_opt	Claude Opus 4.6	0.94 [0.94, 0.95]	0.99 [0.98, 0.99]	0.90 [0.89, 0.91]	0.94 [0.93, 0.95]	0.94 [0.94, 0.95]
	Claude Sonnet 4.6	0.89 [0.88, 0.89]	0.97 [0.97, 0.98]	0.89 [0.88, 0.90]	0.93 [0.93, 0.94]	0.92 [0.92, 0.92]
	Claude Haiku 4.5	0.82 [0.81, 0.82]	0.95 [0.94, 0.97]	0.89 [0.89, 0.90]	0.92 [0.92, 0.93]	0.89 [0.88, 0.89]
	Qwen3-VL	0.89 [0.89, 0.89]	0.98 [0.98, 0.98]	0.87 [0.84, 0.89]	0.92 [0.91, 0.94]	0.92 [0.91, 0.93]
	Qwen3-Next	0.86 [0.86, 0.86]	0.98 [0.98, 0.98]	0.83 [0.83, 0.84]	0.90 [0.90, 0.91]	0.88 [0.87, 0.88]

multipliers are undefined. Measured instead by total cost across the eight claims, prompt caching reduces cost by 2.0x.

We examine the similarity filter’s sensitivity to the threshold. The benchmark’s semantic filters reduce to four distinct predicates, those of C3, C9, C16, and C21; the rest reuse one of these on the same dataset and yield identical curves. Figures 6b and 6c report recall (the fraction of reference-selected tuples above the threshold) and filter rate (the fraction of tuples removed) across thresholds. At the default of 0.15, recall exceeds 0.99 on all four filters with a filter rate of at most 0.02. At higher thresholds, claims with *semantically specific* filter predicates (e.g., C16: “The restaurant review

{text} mentions the number of whiskey varieties available”) maintain high recall, while claims with *semantically broader* predicates (e.g., C21: “The customer support {dialog} contains a customer complaint”) experience a rapid drop in recall. Thus, 0.15 is a conservative default. Integrating more advanced techniques (e.g., learning the threshold [60, 68] or late interaction [43, 66]) is a natural next step.

Four of the six optimizations—relevance sorting, operator fusion, similarity filtering, and prompt caching—maintain verification quality, while disabling early stopping or estimation each lowers precision by roughly 0.06. Both drops trace to C4, a grounded universal claim. With early stopping or estimation enabled, EVERGREEN




Figure 6: Optimization analysis. (a) Cost multipliers (log scale) when each optimization is disabled, relative to evg_opt with Haiku 4.5 (ES: Early Stopping, RS: Relevance Sorting, ECS: Estimation with Confidence Sequences, OF: Operator Fusion, SF: Similarity Filtering). Each point is one claim’s multiplier (mean of three trials); the horizontal line is the geometric mean and the dashed line at $1\times$ marks parity with evg_opt. (b) Recall and (c) filter rate of the similarity filter across thresholds; the dashed line marks the 0.15 default. Points are means over three trials, and shaded regions show the range.

only processes a subset of tuples and correctly finds no counterexample. Without these optimizations, processing all tuples increases the LLM’s likelihood of flagging a false counterexample, flipping the verdict. This effect is not general, however, since estimation can also miss a valid counterexample by sampling. Overall, our optimizations yield substantial cost savings while maintaining quality.

7 RELATED WORK

Section 1 presented work relevant to claim verification. Here, we discuss other related work in more detail.

Semantic query processing engines such as Snowflake’s Cortex AISQL [48], DocETL [68, 81], LOTUS [60], and Palimpzest [49, 64], among many others [4, 16, 18, 21, 23, 24, 32, 47, 51, 65, 75, 78, 82], have emerged in recent years. At their core is a set of semantic operators—including semantic aggregation—which augment traditional database operators with LLMs. To reduce the cost of invoking these operators, novel query optimizers [64, 81] and efficient execution strategies [50, 69, 71, 76, 85, 86, 90] have been proposed, with recent work introducing the SemBench benchmark for evaluating these systems [44]. Yet, the community lacks a systematic approach to verify the correctness of semantic aggregates. Lee et al. [45] describe a vision for semantic integrity constraints as a declarative abstraction for enforcing the correctness of LLM outputs in semantic queries. EVERGREEN is a step towards this vision.

Approximate query processing (AQP) trades accuracy for speed via sampling [1, 14, 33, 59, 92]. Online aggregation [33] introduced progressive refinement of aggregate estimates, allowing users to stop early once estimates appear tight enough; however, continuously monitoring classical CIs invalidates their error guarantees [38]. EVERGREEN’s use of CSs [22, 79] enables valid monitoring without a predetermined sample size. More recent work explores approximating aggregation queries that involve evaluating multimodal semantic operators for filters (e.g., ABAE [39], ThalamusDB [36]) and the statistics of interest (e.g., InQuest [63]). In contrast to these settings, which estimate bounds on the true aggregate values, verification queries only need to resolve boolean comparisons involving the aggregates. This structure enables earlier termination: while traditional AQP relaxes a precise aggregate to an interval estimate, EVERGREEN collapses the interval to a boolean verdict, needing only enough precision to resolve the comparison.

Data provenance originated from the database community [26], with foundational models including data lineage [20], why- and where-provenance [12], and why-not provenance [13]. Seminal work by Green et al. [30] unifies several of these models under provenance semirings, also known as how-provenance. Subsequent work extends semiring semantics to other settings, such as aggregation [3] and full first-order logic with negation [31]. Several database systems support provenance capture and querying [27, 46, 55, 56, 61, 67]; however, none support semantic queries or provenance for first-order logic. In NLP, Zhang et al. [88, 89] formalize provenance for claims as a labeled graph over dependent web sources, tracing how a claim originated and evolved. Unlike this setting, EVERGREEN verifies quantified claims over independent tuples in a relation, motivating its choice of provenance semantics.

8 CONCLUSIONS

We present EVERGREEN, a system that reliably and efficiently verifies claims extracted from semantic aggregates by compiling them into declarative semantic verification queries composed of standard relational and semantic operators. This formulation admits a suite of complementary optimizations that substantially reduce cost while preserving verification quality. Furthermore, each verdict is accompanied by a minimal explanation based on semiring provenance for first-order logic. Looking ahead, a natural extension is to exploit shared computation across the many claims that arise from a semantic aggregate, akin to multi-query optimization in databases. As semantic aggregation becomes a standard feature of production data systems, EVERGREEN demonstrates that the declarative query abstractions already present in these systems can be efficiently leveraged to make their outputs reliable and explainable.

ACKNOWLEDGMENTS

We thank the Snowflake AISQL team for their support and the anonymous reviewers for their constructive feedback. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2439559. 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.

REFERENCES

- [1] Sameer Agarwal, Barzan Mozafari, Aurojit Panda, Henry Milner, Samuel Madden, and Ion Stoica. 2013. BlinkDB: queries with bounded errors and bounded response times on very large data. In *Proceedings of the 8th ACM European Conference on Computer Systems* (Prague, Czech Republic) (*EuroSys '13*). Association for Computing Machinery, New York, NY, USA, 29–42. <https://doi.org/10.1145/2465351.2465355>
- [2] Alibaba Cloud. 2026. Alibaba Cloud Model Studio: Model inference pricing. <https://www.alibabacloud.com/help/en/model-studio/model-pricing>. Accessed: 2026-06-23.
- [3] Yael Amsterdamer, Daniel Deutch, and Val Tannen. 2011. Provenance for aggregate queries. In *Proceedings of the Thirtieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (Athens, Greece) (*PODS '11*). Association for Computing Machinery, New York, NY, USA, 153–164. <https://doi.org/10.1145/1989284.1989302>
- [4] 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*.
- [5] Anthropic. 2025. System Card: Claude Haiku 4.5. <https://www-cdn.anthropic.com/7aad69bf12627d42234e01ee7c36305dc2f6a970.pdf>. Accessed: 2026-03-31.
- [6] Anthropic. 2026. Claude API Docs: Pricing. <https://platform.claude.com/docs/en/about-claude/pricing>. Accessed: 2026-03-31.
- [7] Anthropic. 2026. System Card: Claude Opus 4.6. <https://www-cdn.anthropic.com/0dd865075ad3132672ee0ab40b05a53f14cf5288.pdf>. Accessed: 2026-03-31.
- [8] Anthropic. 2026. System Card: Claude Opus 4.8. <https://www-cdn.anthropic.com/0b4915911bb0d19eca5b5ee635c80ef830a37ea.pdf>. Accessed: 2026-06-23.
- [9] Anthropic. 2026. System Card: Claude Sonnet 4.6. <https://www-cdn.anthropic.com/78073f739564e986ff3e28522761a7a0b4484f84.pdf>. Accessed: 2026-03-31.
- [10] Stuart Axelbrooke. 2017. Customer Support on Twitter. <https://doi.org/10.34740/KAGGLE/DSV/8841>
- [11] Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibao Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. 2025. Qwen3-VL Technical Report. arXiv:2511.21631 [cs.CV] <https://arxiv.org/abs/2511.21631>
- [12] Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. 2001. Why and Where: A Characterization of Data Provenance. In *Database Theory – ICDT 2001*, Jan Van den Bussche and Victor Vianu (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 316–330.
- [13] Adriane Chapman and H. V. Jagadish. 2009. Why not?. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data* (Providence, Rhode Island, USA) (*SIGMOD '09*). Association for Computing Machinery, New York, NY, USA, 523–534. <https://doi.org/10.1145/1559845.1559901>
- [14] Surajit Chaudhuri, Bolin Ding, and Srikanth Kandula. 2017. Approximate Query Processing: No Silver Bullet. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) (*SIGMOD '17*). Association for Computing Machinery, New York, NY, USA, 511–519. <https://doi.org/10.1145/3035918.3056097>
- [15] Atoosa Chegini, Keivan Rezaei, Hamid Eghbalzadeh, and Soheil Feizi. 2025. RePanda: Pandas-powered Tabular Verification and Reasoning. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 32200–32212. <https://doi.org/10.18653/v1/2025.acl-long.1549>
- [16] Shu Chen, Deepti Raghavan, and Ugur Cetintemel. 2025. Continuous Prompts: LLM-Augmented Pipeline Processing over Unstructured Streams. arXiv:2512.03389 [cs.DB] <https://arxiv.org/abs/2512.03389>
- [17] Wenhui Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. TabFact: A Large-scale Dataset for Table-based Fact Verification. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rkeIRhNYDH>
- [18] Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. 2023. Binding Language Models in Symbolic Languages. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=IH1PV42cbF>
- [19] Gordon V. Cormack, Charles L. A. Clarke, and Stefan Buettcher. 2009. Reciprocal rank fusion outperforms condorcet and individual rank learning methods. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Boston, MA, USA) (*SIGIR '09*). Association for Computing Machinery, New York, NY, USA, 758–759. <https://doi.org/10.1145/1571941.1572114>
- [20] Yingwei Cui, Jennifer Widom, and Janet L. Wiener. 2000. Tracing the lineage of view data in a warehousing environment. *ACM Trans. Database Syst.* 25, 2 (June 2000), 179–227. <https://doi.org/10.1145/357775.357777>
- [21] 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 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=7SGOv5W5z>
- [22] D. A. Darling and Herbert Robbins. 1967. CONFIDENCE SEQUENCES FOR MEAN, VARIANCE, AND MEDIAN. *Proceedings of the National Academy of Sciences* 58, 1 (1967), 66–68. <https://doi.org/10.1073/pnas.58.1.66>
- [23] Uélison Jean Lopes dos Santos, Alessandro Ferri, Szilard Nistor, Riccardo Tommasini, Carsten Binnig, and Manisha Luthra. 2025. Towards a Multimodal Stream Processing System. arXiv:2510.14631 [cs.DB] <https://arxiv.org/abs/2510.14631>
- [24] Till Döhmen. 2024. INTRODUCING THE PROMPT() FUNCTION: USE THE POWER OF LLMs WITH SQL! <https://motherduck.com/blog/sql-llm-prompt-function-gpt-models/>. Accessed: 2026-03-31.
- [25] Guy Feigenblat, Chulaka Gunasekara, Benjamin Sznajder, Sachindra Joshi, David Konopnicki, and Ranit Aharonov. 2021. TWEETSUMM - A Dialog Summarization Dataset for Customer Service. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 245–260. <https://doi.org/10.18653/v1/2021.findings-emnlp.24>
- [26] Boris Glavic. 2021. Data Provenance. *Found. Trends Databases* 9, 3-4 (April 2021), 209–441. <https://doi.org/10.1561/19000000068>
- [27] Boris Glavic and Gustavo Alonso. 2009. Perm: Processing Provenance and Data on the Same Data Model through Query Rewriting. In *Proceedings of the 2009 IEEE International Conference on Data Engineering (ICDE '09)*. IEEE Computer Society, USA, 174–185. <https://doi.org/10.1109/ICDE.2009.15>
- [28] Google. 2026. Gemini 3.5 Flash Model Card. <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-3-5-Flash-Model-Card.pdf>. Accessed: 2026-06-23.
- [29] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, et al. 2024. The Llama 3 Herd of Models. arXiv:2407.21783 [cs.AI] <https://arxiv.org/abs/2407.21783>
- [30] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (Beijing, China) (*PODS '07*). Association for Computing Machinery, New York, NY, USA, 31–40. <https://doi.org/10.1145/1265530.1265535>
- [31] Erich Grädel and Val Tannen. 2024. Provenance Analysis and Semiring Semantics for First-Order Logic. arXiv:2412.07986 [cs.LO] <https://arxiv.org/abs/2412.07986>
- [32] Jian He and Vaibhav Sethi. 2025. Announcing BigQuery-managed AI functions for better SQL. <https://cloud.google.com/blog/products/data-analytics/sql-reimagined-for-the-ai-era-with-bigquery-ai-functions>. Accessed: 2026-03-31.
- [33] Joseph M. Hellerstein, Peter J. Haas, and Helen J. Wang. 1997. Online aggregation. *SIGMOD Rec.* 26, 2 (June 1997), 171–182. <https://doi.org/10.1145/253262.253291>
- [34] Qisheng Hu, Quanyu Long, and Wenya Wang. 2025. BOOST: Bootstrapping Strategy-Driven Reasoning Programs for Program-Guided Fact-Checking. arXiv:2504.02467 [cs.AI] <https://arxiv.org/abs/2504.02467>
- [35] Tharushi Jayasekara and Immanuel Trummer. 2025. CEDAR: A System for Cost-Efficient Data-Driven Claim Verification. *Proc. VLDB Endow.* 18, 11 (July 2025), 4492–4504. <https://doi.org/10.14778/3749646.3749708>
- [36] Saehan Jo and Immanuel Trummer. 2024. ThalamusDB: Approximate Query Processing on Multi-Modal Data. *Proc. ACM Manag. Data* 2, 3, Article 186 (May 2024), 26 pages. <https://doi.org/10.1145/3654989>
- [37] Saehan Jo, Immanuel Trummer, Weicheng Yu, Xuezhi Wang, Cong Yu, Daniel Liu, and Niyati Mehta. 2019. Verifying Text Summaries of Relational Data Sets. In *Proceedings of the 2019 International Conference on Management of Data* (Amsterdam, Netherlands) (*SIGMOD '19*). Association for Computing Machinery, New York, NY, USA, 299–316. <https://doi.org/10.1145/3299869.3300074>
- [38] Ramesh Johari, Pete Koomen, Leonid Pekelis, and David Walsh. 2017. Peeking at A/B Tests: Why it matters, and what to do about it. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Halifax, NS, Canada) (*KDD '17*). Association for Computing Machinery, New York, NY, USA, 1517–1525. <https://doi.org/10.1145/3097983.3097992>
- [39] Daniel Kang, John Guibas, Peter Bailis, Tatsunori Hashimoto, Yi Sun, and Matei Zaharia. 2021. Accelerating approximate aggregation queries with expensive predicates. *Proc. VLDB Endow.* 14, 11 (July 2021), 2341–2354. <https://doi.org/10.14778/3476249.3476285>
- [40] Georgios Karagiannis, Mohammed Saeed, Paolo Papotti, and Immanuel Trummer. 2020. Scrutinizer: a mixed-initiative approach to large-scale, data-driven claim verification. *Proc. VLDB Endow.* 13, 12 (July 2020), 2508–2521. <https://doi.org/10.14778/3407790.3407841>

- [41] Omar Khattab, Christopher Potts, and Matei Zaharia. 2021. Baleen: Robust Multi-Hop Reasoning at Scale via Condensed Retrieval. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 27670–27682. https://proceedings.neurips.cc/paper_files/paper/2021/file/e8b1cbd05f6e6a358a81dee52493dd06-Paper.pdf
- [42] Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan A. Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. DSPy: Compiling Declarative Language Model Calls into State-of-the-Art Pipelines. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=sY5N0zY5Od>
- [43] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 39–48. <https://doi.org/10.1145/3397271.3401075>
- [44] Jiale Lao, Andreas Zimmerer, Olga Ovcharenko, Tianji Cong, Matthew Russo, Gerardo Vitagliano, Michael Cochez, Fatma Özcan, Gautam Gupta, Thibaud Hottelier, H. V. Jagadish, Kris Kissel, Sebastian Schelter, Andreas Kipf, and Immanuel Trummer. 2026. SemBench: A Benchmark for Semantic Query Processing Engines. *Proc. VLDB Endow.* 19, 8 (July 2026), 1754–1767. <https://doi.org/10.14778/3811243.3811249>
- [45] Alexander W. Lee, Justin Chan, Michael Fu, Nicolas Kim, Akshay Mehta, Deepti Raghavan, and Uğur Çetintemel. 2025. Semantic Integrity Constraints: Declarative Guardrails for AI-Augmented Data Processing Systems. *Proc. VLDB Endow.* 18, 11 (July 2025), 4073–4080. <https://doi.org/10.14778/3749646.3749677>
- [46] Seokki Lee, Bertram Ludäscher, and Boris Glavic. 2019. PUG: a framework and practical implementation for why and why-not provenance. *The VLDB Journal* 28, 1 (Feb. 2019), 47–71. <https://doi.org/10.1007/s00778-018-0518-5>
- [47] Zequn Li, Yuanhao Zhong, Chengliang Chai, Zhaoze Sun, Yuhao Deng, Ye Yuan, Guoren Wang, and Lei Cao. 2025. DocDB: A Database for Unstructured Document Analysis. *Proc. VLDB Endow.* 18, 12 (Aug. 2025), 5387–5390. <https://doi.org/10.14778/3750601.3750678>
- [48] Paweł Liskowski, Benjamin Han, Paritosh Aggarwal, Bowei Chen, Boxin Jiang, Nitish Jindal, Zihan Li, Aaron Lin, Kyle Schmaus, Jay Tayade, Weicheng Zhao, Anupam Datta, Nathan Wiegand, and Dimitris Tsirigiannis. 2026. Cortex AISQL: A Production SQL Engine for Unstructured Data. In *Companion of the International Conference on Management of Data (India) (SIGMOD Companion '26)*. Association for Computing Machinery, New York, NY, USA, 400–412. <https://doi.org/10.1145/3788853.3803093>
- [49] 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*.
- [50] Shu Liu, Asim Biswal, Amog Kamsetty, Audrey Cheng, Luis Gaspar Schroeder, Liana Patel, Shiyi Cao, Xiangxi Mo, Ion Stoica, Joseph E. Gonzalez, and Matei Zaharia. 2025. Optimizing LLM Queries in Relational Data Analytics Workloads. In *Eighth Conference on Machine Learning and Systems*. <https://openreview.net/forum?id=R7bK9yycHp>
- [51] Duo Lu, Siming Feng, Jonathan Zhou, Franco Solleza, Malte Schwarzkopf, and Uğur Çetintemel. 2025. VectraFlow: Integrating Vectors into Stream Processing. In *CIDR*.
- [52] Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. 2003. TAG: a Tiny AGgregation service for ad-hoc sensor networks. *SIGOPS Oper. Syst. Rev.* 36, SI (Dec. 2003), 131–146. <https://doi.org/10.1145/844128.844142>
- [53] Luke Merrick, Danmei Xu, Gaurav Nuti, and Daniel Campos. 2024. Arctic-Embed: Scalable, Efficient, and Accurate Text Embedding Models. arXiv:2405.05374 [cs.CL] <https://arxiv.org/abs/2405.05374>
- [54] Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FactScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 12076–12100. <https://doi.org/10.18653/v1/2023.emnlp-main.741>
- [55] Haneen Mohammed, Charlie Summers, Sugosh Kaushik, and Eugene Wu. 2023. SmokedDuck Demonstration: SQLStepper. In *Companion of the 2023 International Conference on Management of Data (Seattle, WA, USA) (SIGMOD '23)*. Association for Computing Machinery, New York, NY, USA, 183–186. <https://doi.org/10.1145/3555041.3589731>
- [56] Haneen Mohammed, Alexander Yao, Charlie Summers, Hongbin Zhong, Gromit Yeuk-Yin Chan, Subrata Mitra, Lampros Flokas, and Eugene Wu. 2024. FaDE: More Than a Million What-Ifs Per Second. *Proc. VLDB Endow.* 18, 4 (Dec. 2024), 943–955. <https://doi.org/10.14778/3717755.3717757>
- [57] OpenAI. 2026. GPT-5.5 System Card. <https://deploymentsafety.openai.com/gpt-5-5/gpt-5-5.pdf>. Accessed: 2026-06-23.
- [58] Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan, and Preslav Nakov. 2023. Fact-Checking Complex Claims with Program-Guided Reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 6981–7004. <https://doi.org/10.18653/v1/2023.acl-long.386>
- [59] Yongjoo Park, Barzan Mozafari, Joseph Sorenson, and Junhao Wang. 2018. VerdictDB: Universalizing Approximate Query Processing. In *Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD '18)*. Association for Computing Machinery, New York, NY, USA, 1461–1476. <https://doi.org/10.1145/3183713.3196905>
- [60] Liana Patel, Siddharth Jha, Melissa Pan, Harshit Gupta, Parth Asawa, Carlos Guestrin, and Matei Zaharia. 2025. Semantic Operators and Their Optimization: Enabling LLM-Based Data Processing with Accuracy Guarantees in LOTUS. *Proc. VLDB Endow.* 18, 11 (July 2025), 4171–4184. <https://doi.org/10.14778/3749646.3749685>
- [61] Fotis Psallidas and Eugene Wu. 2018. Smoke: fine-grained lineage at interactive speed. *Proc. VLDB Endow.* 11, 6 (Feb. 2018), 719–732. <https://doi.org/10.14778/3199517.3199522>
- [62] QwenTeam. 2025. Qwen3-Next: Towards Ultimate Training & Inference Efficiency. <https://qwen.ai/blog?id=qwen3-next>. Accessed: 2026-06-23.
- [63] Matthew Russo, Tatsunori Hashimoto, Daniel Kang, Yi Sun, and Matei Zaharia. 2023. Accelerating Aggregation Queries on Unstructured Streams of Data. *Proc. VLDB Endow.* 16, 11 (July 2023), 2897–2910. <https://doi.org/10.14778/3611479.3611496>
- [64] Matthew Russo, Chunwei Liu, Sivaprasad Sudhir, Gerardo Vitagliano, Michael Cafarella, Tim Kraska, and Samuel Madden. 2026. Abacus: A Cost-Based Optimizer for Semantic Operator Systems. *Proc. VLDB Endow.* 19, 5 (May 2026), 1060–1073. <https://doi.org/10.14778/3796195.3796215>
- [65] Gabriele Sanmartino, Matthias Urban, Paolo Papotti, and Carsten Binnig. 2026. The Stretto Execution Engine for LLM-Augmented Data Systems. arXiv:2602.04430 [cs.DB] <https://arxiv.org/abs/2602.04430>
- [66] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 3715–3734. <https://doi.org/10.18653/v1/2022.naacl-main.272>
- [67] Aryak Sen, Silviu Maniu, and Pierre Senellart. 2025. ProvSQL: A General System for Keeping Track of the Provenance and Probability of Data. arXiv:2504.12058 [cs.DB] <https://arxiv.org/abs/2504.12058>
- [68] Shreya Shankar, Tristan Chambers, Tarak Shah, Aditya G. Parameswaran, and Eugene Wu. 2025. DocETL: Agentic Query Rewriting and Evaluation for Complex Document Processing. *Proc. VLDB Endow.* 18, 9 (May 2025), 3035–3048. <https://doi.org/10.14778/3746405.3746426>
- [69] Shreya Shankar, Sepanta Zeighami, and Aditya Parameswaran. 2026. Task Cascades for Efficient Unstructured Data Processing. *Proc. ACM Manag. Data* 4, 1, Article 88 (April 2026), 26 pages. <https://doi.org/10.1145/3786702>
- [70] Snowflake. n.d. Snowflake Cortex AI. <https://www.snowflake.com/en/product/features/cortex/>. Accessed: 2026-03-31.
- [71] Zhaoze Sun, Chengliang Chai, Qiyang Deng, Kaisen Jin, Xinyu Guo, Han Han, Ye Yuan, Guoren Wang, and Lei Cao. 2025. QUEST: Query Optimization in Unstructured Document Analysis. *Proc. VLDB Endow.* 18, 11 (July 2025), 4560–4573. <https://doi.org/10.14778/3749646.3749713>
- [72] 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>
- [73] Michael Theologitis, Preetam Prabhu Srikar Dammu, Chirag Shah, and Dan Suciu. 2026. ClaimDB: A Fact Verification Benchmark over Large Structured Data. arXiv:2601.14698 [cs.CL] <https://arxiv.org/abs/2601.14698>
- [74] Michael Theologitis and Dan Suciu. 2026. Thucy: An LLM-based Multi-Agent System for Claim Verification across Relational Databases. arXiv:2512.03278 [cs.DB] <https://arxiv.org/abs/2512.03278>
- [75] Matthias Urban and Carsten Binnig. 2023. CAESURA: Language Models as Multi-Modal Query Planners. In *CIDR*.
- [76] Matthias Urban and Carsten Binnig. 2024. ELEET: Efficient Learned Query Execution over Text and Tables. *Proc. VLDB Endow.* 17, 13 (Sept. 2024), 4867–4880. <https://doi.org/10.14778/3704965.3704989>
- [77] Haoran Wang and Kai Shu. 2023. Explainable Claim Verification via Knowledge-Grounded Reasoning with Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 6288–6304. <https://doi.org/10.18653/v1/2023.findings-emnlp.416>

- [78] Jiayi Wang and Guoliang Li. 2025. AOP: Automated and Interactive LLM Pipeline Orchestration for Answering Complex Queries. In *CIDR*.
- [79] Ian Waudby-Smith and Aaditya Ramdas. 2023. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 86, 1 (02 2023), 1–27. <https://doi.org/10.1093/jrsssb/qkad009> arXiv:<https://academic.oup.com/jrsssb/article-pdf/86/1/1/56961777/qkad009.pdf>
- [80] Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. Long-form factuality in large language models. 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., 80756–80827. <https://doi.org/10.52202/079017-2567>
- [81] Lindsey Linxi Wei, Shreya Shankar, Sepanta Zeighami, Yeounoh Chung, Fatma Ozcan, and Aditya G. Parameswaran. 2026. Multi-Objective Agentic Rewrites for Unstructured Data Processing. arXiv:2512.02289 [cs.DB] <https://arxiv.org/abs/2512.02289>
- [82] Patrick Wendell, Eric Peter, Nicolas Pelaez, Jianwei Xie, Vinny Vijayakumar, Linhong Liu, and Shitao Li. 2023. Introducing AI Functions: Integrating Large Language Models with Databricks SQL. <https://www.databricks.com/blog/2023/04/18/introducing-ai-functions-integrating-large-language-models-databricks-sql.html>. Accessed: 2026-03-31.
- [83] Zhuohan Xie, Rui Xing, Yuxia Wang, Jiahui Geng, Hasan Iqbal, Dhruv Sahnan, Iryna Gurevych, and Preslav Nakov. 2025. FIRE: Fact-checking with Iterative Retrieval and Verification. In *Findings of the Association for Computational Linguistics: NAACL 2025*, Luis Chiruzzo, Alan Ritter, and Lu Wang (Eds.). Association for Computational Linguistics, Albuquerque, New Mexico, 2901–2914. <https://doi.org/10.18653/v1/2025.findings-naacl.158>
- [84] Yelp. 2023. Yelp Open Dataset. <https://business.yelp.com/data/resources/open-dataset/>. Accessed: 2026-03-31.
- [85] Sepanta Zeighami, Shreya Shankar, and Aditya Parameswaran. 2025. Cut Costs, Not Accuracy: LLM-Powered Data Processing with Guarantees. *Proc. ACM Manag. Data* 3, 6, Article 311 (Dec. 2025), 26 pages. <https://doi.org/10.1145/3769776>
- [86] Sepanta Zeighami, Shreya Shankar, and Aditya Parameswaran. 2025. Featurized-Decomposition Join: Low-Cost Semantic Joins with Guarantees. arXiv:2512.05399 [cs.DB] <https://arxiv.org/abs/2512.05399>
- [87] Alex L. Zhang, Tim Kraska, and Omar Khattab. 2026. Recursive Language Models. arXiv:2512.24601 [cs.AI] <https://arxiv.org/abs/2512.24601>
- [88] Yi Zhang, Zachary Ives, and Dan Roth. 2020. “Who said it, and Why?” Provenance for Natural Language Claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 4416–4426. <https://doi.org/10.18653/v1/2020.acl-main.406>
- [89] Yi Zhang, Zachary Ives, and Dan Roth. 2021. What is Your Article Based On? Inferring Fine-grained Provenance. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online. <https://doi.org/10.18653/v1/2021.acl-long.458>
- [90] Fuheng Zhao, Jiayue Chen, Yiming Pan, Tahseen Rabbani, Sohaib, Divyakant Agrawal, Amr El Abbadi, Paritosh Aggarwal, Anupam Datta, and Dimitris Tsirigiannis. 2025. Access Paths for Efficient Ordering with Large Language Models. arXiv:2509.00303 [cs.DB] <https://arxiv.org/abs/2509.00303>
- [91] 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.
- [92] Yuxuan Zhu, Tengjun Jin, Stefanos Baziotis, Chengsong Zhang, Charith Mendis, and Daniel Kang. 2025. PilotDB: Database-Agnostic Online Approximate Query Processing with A Priori Error Guarantees. *Proc. ACM Manag. Data* 3, 3, Article 198 (June 2025), 28 pages. <https://doi.org/10.1145/3725335>