STEER: Reducing Inference Cost in Relational Foundation Models through Semantically Informed Sampling

Abdalla Mohamed
University of Texas at Arlington
axm0993@mavs.uta.edu

Ashraf Aboulnaga
University of Texas at Arlington
ashraf.aboulnaga@uta.edu

ABSTRACT

Relational foundation models (RFMs) are pretrained once on a collection of relational databases and prediction tasks, and then applied zero-shot to previously unseen databases and tasks. To make a prediction for a target row, an RFM samples a neighborhood of rows linked to that row through foreign keys and uses this neighborhood as its inference context. Lowering inference cost is an important goal for any foundation model, and for RFMs this cost grows with the size of the context. The simplest ways to shrink the context is to drop some of the sampled rows, but this ignores the semantics of the database schema, so it is as likely to discard informative rows as uninformative ones. We propose STEER, a sampling approach that shrinks the inference context by concentrating it on the tables most relevant to the prediction task at hand. STEER obtains relevance information by prompting a large language model to rank the foreign-key edges of the database schema into relevance tiers for the given task, and then maps each tier to a probability of following that edge during traversal. Because the ranking uses only the schema, it is computed once per task and reused across all subsequent predictions, amortizing its cost. We evaluate STEER on three state-of-the-art RFMs (RT, RT-J, and Griffin) and show that it reduces inference context size by about 40% on average while maintaining, and in some cases improving, prediction accuracy.

Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/lids-lab/steer.

1 INTRODUCTION

Relational deep learning (RDL) aims to train machine learning models directly on the structure of relational databases, without flattening the database into a single table and without the manual, dataset-specific feature engineering required by traditional tabular models [4, 32]. In RDL, a relational database is represented as a graph in which each row of each table is a node, and each foreign-key reference is an edge connecting two nodes [8]. A model can then propagate information along these edges, allowing the representation learned for a given row to be enriched with information from related rows. Early RDL models were shown to be effective across a wide range of classification and regression tasks, but they were trained for a specific database and task, requiring retraining from scratch for each new database or task [5, 7].

More recent work in RDL has focused on *relational foundation models (RFMs)*, which are pretrained once on a collection of databases and can subsequently make zero-shot predictions for tasks on previously unseen databases, without updating their model weights. An early example is Griffin [33], which is built around a graph neural network (GNN) architecture. Other examples include the closed-source, closed-weight, commercially available KumoRFM [9] and KumoRFM-2 [15] models. Another prominent example is the open-source Relational Transformer (RT) [29], which is built around a transformer architecture and has demonstrated strong performance across a range of evaluations. A recent follow-up, RT-J [30], retains the RT architecture while introducing a large-scale pretraining recipe aimed at label-efficient prediction.

At inference time, an RFM must make predictions on databases it has never seen before and for which it has no prior knowledge of the structure or contents. Consequently, for each prediction task, the RFM constructs a *sample* of the target database containing the information needed to make the prediction. The model sees the database only through the lens of this sample, so the sampler effectively determines what information is available to the model at inference time. This sample therefore serves as the *context* for the prediction, and the cost of inference depends strongly on the size and complexity of this context.

Reducing inference cost is essential to the practical deployment of foundation models of all types, including RFMs. Unlike training, which is performed once, inference incurs a recurring cost for every prediction request, so even modest savings are multiplied across users, applications, and the lifetime of a deployed model. Accordingly, considerable attention in both academia and industry has focused on reducing the inference cost of large language models (LLMs), which are by far the dominant class of foundation models and the focus of most current AI activity [40]. Commercial LLM inference is typically priced by the number of input and output tokens processed, making context size not only a computational concern for the model provider but also a direct contributor to the monetary cost paid by model users. Approaches to reducing LLM inference cost include specialized inference hardware, such as Groq [1] and AWS Inferentia [3], as well as algorithmic techniques such as quantization [37] and speculative decoding [20]. Another important line of work reduces inference cost by reducing the amount of context processed by the model, for example through context compression [22] or selective attention [14]. In this paper, we pursue this latter direction for RFMs by constructing smaller contexts through semantically informed database sampling.

To motivate our sampling approach, we first describe how current RFM samplers construct inference context. An RFM prediction pertains to a particular row in the database, which we term the *seed row*. Because RDL represents each row as a node in a graph, the seed row corresponds to a *seed node* in the graph. RFM samplers construct the prediction context by sampling a neighborhood around this seed node. Starting from the seed node, the sampler traverses the graph along edges representing foreign-key links and

adds the rows encountered during traversal to the sample. The process continues for a number of steps until a stopping condition is met, with the assumption that the resulting neighborhood provides informative context for prediction.

The cost of RFM inference grows with the size of the sampled context, since the computational cost of the model's forward pass increases with the amount of input it must process. Transformer-based RFMs consume the sample as a sequence of tokens, whereas GNN-based RFMs operate directly on the sampled subgraph. In either case, smaller samples reduce inference cost. A simple way to shrink the sample is to drop sampled rows at random. However, random dropping ignores the relevance of individual rows to the prediction task and can therefore substantially degrade prediction accuracy when informative rows are dropped.

Consider, for example, the Formula 1 database in the RelBench benchmark [31], the standard benchmark for relational deep learning. This database includes tables such as `drivers`, `results`, `standings`, and `qualifying`. Suppose the task is to predict whether a given driver (represented by a row in the `drivers` table) will fail to finish an upcoming race. The `results` table records each driver's previous race outcomes, including past instances of failing to finish, and therefore provides a more direct signal for this task than `standings` or `qualifying`, which describe a driver's ranking and grid position. A sampler that drops rows at random is just as likely to discard informative `results` rows as less useful `standings` or `qualifying` rows, potentially filling its limited context with information that contributes little to the prediction. A better approach, and the one we pursue in this paper, is to account for the expected usefulness of each table to the prediction task when deciding what information to sample.

Our approach is based on the observation that an important source of knowledge about the likely usefulness of tables is available without examining a single row of the database: the semantic structure of the schema itself. Table and column names are chosen to describe the real-world concepts they represent, and these descriptions can often indicate whether a table is likely to be relevant to a particular prediction task before any of its values are inspected. Building on this observation, we propose *STEER (Schema-guided, Token-Efficient Edge Ranking)*, a simple yet effective sampling approach that injects schema-derived semantic information about the usefulness of different tables for the prediction task into the sampling process. STEER replaces random row selection with a structured relevance ranking that governs the graph traversal itself. Rather than sampling broadly and then randomly dropping rows, STEER reduces sample size by actively choosing which edges to traverse and, consequently, which tables to sample.

Another important observation underpinning STEER is that modern LLMs are a good source of deep semantic information about database schemas and how their table names, column names, and foreign-key connectivity relate to prediction tasks. Therefore, STEER uses an LLM as a source of semantic knowledge and prompts this LLM to rank every edge leaving each table, in both traversal directions, into one of four tiers based on the relevance of this edge to the specific prediction task. Each tier is assigned a traversal probability, which the sampler uses to decide whether to follow an edge, rather than unconditionally traversing every edge, as in the default sampler. We also present an extension of STEER that identifies task-relevant columns, although the primary focus of STEER is on controlling graph traversal.

We implement STEER on the RT, RT-J, and Griffin RFMs, modifying each RFM's sampler, and we evaluate it on the RelBench benchmark, using tiered rankings generated by the Claude and GPT frontier LLMs. Across this evaluation, on average, STEER reduces sample size by roughly 40% while matching, and sometimes slightly improving, the prediction accuracy of the baseline sampler operating with the full sample size. In contrast, random row dropping can (by construction) reduce sample size as much as needed, but this reduction comes at a steep cost in prediction accuracy. In summary, our contributions are as follows:

- A schema-aware LLM prompting strategy that produces tiered semantic relevance rankings for graph-traversal edges with respect to a given prediction task.
- A probabilistic sampling mechanism that uses these rankings to steer inference-time traversal, reducing sample size without sacrificing
  prediction accuracy.
- An optional column-level filtering mechanism that aims to remove columns that are unlikely to be informative for the prediction task.
- An empirical evaluation across three relational foundation models (RT, RT-J, and Griffin) and two frontier LLMs (Claude and GPT), showing that, on average, STEER reduces sample size by roughly 40% while matching (or slightly improving) the prediction accuracy of baseline sampling with a full-size context.

2 BACKGROUND: SAMPLING IN RFMS

In this section, we provide a more detailed description of the sampling process used by RFMs. STEER modifies this process by incorporating semantic information into the sampling decisions.

Most predictive tasks of interest do not ask questions that a database already answers directly, but instead concern outcomes that have not yet occurred, such as whether a customer will churn next month or whether a Formula 1 driver will finish an upcoming race. To express such a prediction task to an RFM, a *task table* [29] is introduced alongside the existing database tables. Each row of the task table specifies (1) the entity for which a prediction is to be made (e.g., a customer or driver), specified through a foreign key pointing to that entity's row in the database; (2) a timestamp indicating the point in time at which the prediction is made; and (3) a label column containing the outcome to be predicted. In a real deployment setting, the label value is not yet known and is therefore left blank. In an evaluation setting, the label is known but hidden from the model at prediction time and revealed only after the prediction is made, so it can be used for evaluation.

Different prediction tasks require different task tables. For example, predicting whether a customer will churn requires a different task table from predicting whether a Formula 1 driver will finish a race. A task table contains one or more rows corresponding to entities for which predictions are required, that is, rows whose label values are unknown or hidden from the model. It may also contain rows corresponding to *example entities*, for which the prediction outcome is known and exposed to the RFM. These examples provide the model with information about the nature of the prediction task and play an important role in improving prediction accuracy.

The task table is linked to the database through the same graph representation used by RDL (every row is a node and every foreign-key reference is an edge). Because each task-table row contains a foreign key linking it to an entity in the database, the task-table row can itself be treated as another node in the graph, connected to the rest of the database through this foreign-key edge.

An alternative to explicitly constructing a task table is to present the prediction task through a declarative query language, such as the Predictive Query Language (PQL) [17]. PQL allows a user to express a prediction request as a query, much as SQL expresses a data-retrieval request. Such a query is ultimately compiled into the same underlying specification of an entity, timestamp, and prediction target that a task-table row encodes. We therefore describe RFM inference in terms of task tables throughout this paper.

To make a prediction, an RFM takes a task-table row as the starting seed row and samples a portion of the surrounding database graph by traversing foreign-key links and adding encountered rows to the sample. Using the full graph for prediction is generally infeasible because the database graph may be too large to process in its entirety. The sampled portion of the graph is therefore provided to the model as input, and the model produces a prediction for the label associated with the seed row.

The traversal distinguishes two directions in which neighboring rows can be reached. First, it can move from a foreign-key *child* row to the primary-key *parent* row that it references. For example, the traversal may move from a `results` row to the corresponding `drivers` row. We refer to this as an F→P link. Because a table has a fixed, typically small, number of foreign-key columns, and each foreign-key value references at most one parent row, F→P traversal has bounded fan-out.

Second, the traversal can move in the opposite direction, from a primary-key parent row to the foreign-key child rows that reference it. For example, it may move from a `drivers` row to all `results` rows that reference that driver. We refer to this as a P→F link. This direction can have very high fan-out because a single parent row may be referenced by many child rows. Consequently, RFM samplers typically limit how many P→F neighbors are explored and select among the available candidates randomly. As a result, which particular rows ultimately appear in the sample is determined by chance rather than by their expected relevance to the prediction task. STEER modifies this by using semantic information to focus traversal on edges that are more likely to be useful.

The precise traversal strategy varies across RFMs, but all construct a bounded sample around the seed row and exclude *temporally invalid* rows, that is, rows whose timestamps are later than the seed row's timestamp. This prevents *temporal leakage*. Griffin [33], which is GNN-based, samples a fixed number B of neighbors per node for K hops around the seed (typical values are $B = 20$ and $K = 2$). RT [29], which is transformer-based, constructs a context containing at most L tokens using a modified breadth-first traversal over foreign-key links. Starting from the seed row, RT visits rows, adds their non-missing feature values as tokens to the context, and places neighboring rows on a *frontier* for subsequent exploration. RT always includes F→P parent rows and subsamples P→F child rows subject to a width bound w, continuing until the context budget is exhausted or no additional rows remain reachable. RT-J [30] retains RT's cell-token representation but replaces its basic

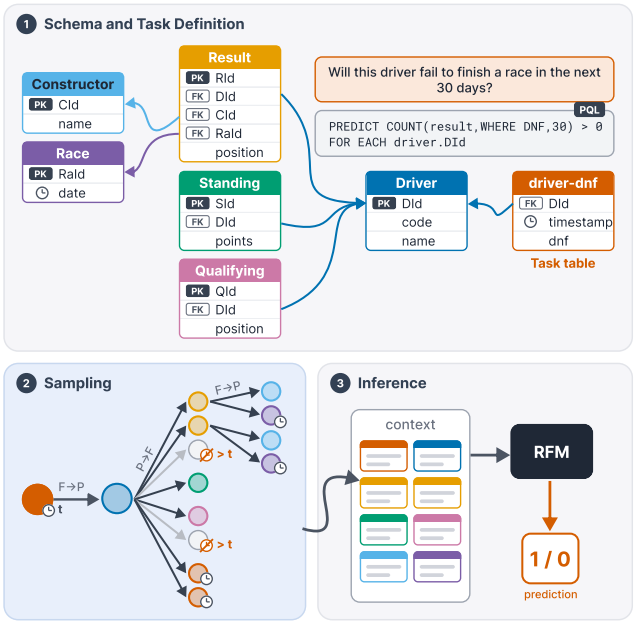


Figure 1: RFM inference on a relational database. (1) The schema, with primary (PK) and foreign keys (FK), and a prediction task defined in English, in PQL, and as a task table whose foreign key references the entity being predicted for. (2) Starting from the task row (seed), the sampler follows F→P and P→F links to collect related rows (node colors match the tables in (1)); rows dated after the prediction timestamp t are excluded. (3) The sampled rows form the context given to the RFM, which outputs the prediction.

neighborhood expansion with a retrieval-based sampling strategy. It performs multiple short random walks from the seed row, uses visitation frequency to identify promising rows, and then applies a local search to fill a bounded context window. STEER is designed to operate with any of these samplers because its decisions are made at the level of which foreign-key edge to follow while sampling.

Figure 1 illustrates the end-to-end RFM inference process on a small example database. The figure shows the database schema and task table, the portion of the graph sampled for a particular prediction, and the resulting model output.

3 CONTROLLING THE TRAVERSAL

As discussed in Section 2, an RFM constructs its sample for inference by traversing the database graph outward from the seed row. STEER aims to control this graph traversal, and to steer it towards more useful and informative rows. We start by presenting the *control points* at which STEER intervenes in the traversal.

The details of the traversal vary between RFMs, but all RFMs follow the overall abstract process described in Algorithm 1. The blue lines in the algorithm are added by STEER and should be ignored for now. The steps Select, Parents, Children, and Done are abstract functions that are instantiated differently for different RFMs, and define the RFM's specific sampling algorithm. The sampler maintains a frontier F of rows to visit. At each step in the sampling loop, the sampler chooses one row r from this frontier,

Algorithm 1 Adding STEER to an RFM sampler. SELECT, PARENTS, CHILDREN, and DONE are abstract functions whose implementation defines the RFM's sampling steps. These functions are left unchanged by STEER. The blue lines are STEER's only additions.

Require: seed row s; row timestamps $\mathcal{T}$; probabilities $\Phi_{T,d}(e) \in [0,1]$ for prediction task T, direction d, and schema edge e
Ensure: sample S

1: $S \leftarrow \{\}$; $F \leftarrow \{s\}$ ▷ sample, frontier
2: **while** $F \neq \{\}$ **and not** DONE(S) **do**
3: $r \leftarrow$ SELECT(F); $F \leftarrow F \setminus \{r\}$
4: $S \leftarrow S \cup \{r\}$
5: $N_{\text{FP}} \leftarrow$ PARENTS(r)
6: $N_{\text{FP}} \leftarrow$ STEERFILTER(r, N_{FP}, F→P) ▷ filter parents
7: $N_{\text{PF}} \leftarrow \{q \in$ CHILDREN(r) $\mid \mathcal{T}(q) \leq \mathcal{T}(s)\}$ ▷ no time leakage
8: $N_{\text{PF}} \leftarrow$ STEERFILTER(r, N_{PF}, P→F) ▷ filter children
9: $F \leftarrow F \cup N_{\text{FP}} \cup N_{\text{PF}}$ ▷ add to frontier
10: **return** S

11: **function** STEERFILTER(r, N, d)
12: **return** $\{ q \in N \mid \text{Uniform}(0,1) < \Phi_{T,d}(\text{edge}(r,q)) \}$

Algorithm 2 Adding STEER to the RT sampler [29], an instance of Algorithm 1. Note that the blue lines added by STEER are unchanged from Algorithm 1. The way that STEER is added to a sampler does not change with the specifics of that sampler.

Require: seed row s; row timestamps $\mathcal{T}$; context length L; width bound w; probabilities $\Phi_{T,d}(e) \in [0,1]$ for prediction task T, direction d, and schema edge e
Ensure: set of cells C (the context)

1: $C \leftarrow \{\}$; $F \leftarrow \{s\}$; $V \leftarrow \{\}$ ▷ context, frontier, visited
2: **while** $F \neq \{\}$ **and** $|C| < L$ **do** ▷ DONE: context length L reached
3: $r \leftarrow$ SELECT(F); $F \leftarrow F \setminus \{r\}$
4: **if** $r \in V$ **then**
5: **continue**
6: $V \leftarrow V \cup \{r\}$; $C \leftarrow C \cup \mathcal{C}(r)$ ▷ add r's non-missing cells
7: $N_{\text{FP}} \leftarrow \mathcal{N}_{\text{F}\to\text{P}}(r)$ ▷ PARENTS: all parents of r
8: $N_{\text{FP}} \leftarrow$ STEERFILTER(r, N_{FP}, F→P) ▷ filter parents
9: $N_{\text{PF}} \leftarrow \{q \in \mathcal{N}_{\text{P}\to\text{F}}(r) \mid \mathcal{T}(q) \leq \mathcal{T}(s)\}$ ▷ no time leakage
10: $N_{\text{PF}} \leftarrow$ RANDOMSAMPLE(N_{PF}, w) ▷ CHILDREN: at most w
11: $N_{\text{PF}} \leftarrow$ STEERFILTER(r, N_{PF}, P→F) ▷ filter children
12: $F \leftarrow F \cup N_{\text{FP}} \cup N_{\text{PF}}$ ▷ add to frontier
13: **return** C

14: **function** SELECT(F) ▷ RT's SELECT
15: $R \leftarrow \{r \in F \mid r$ was added via an F→P link$\}$
16: **if** $R = \{\}$ **then**
17: $R \leftarrow \arg\min_{r \in F}$ HopDistance(r, s) ▷ rows closest to s
18: **return** RANDOMSAMPLE(R, 1)

19: **function** STEERFILTER(r, N, d)
20: **return** $\{ q \in N \mid \text{Uniform}(0,1) < \Phi_{T,d}(\text{edge}(r,q)) \}$

determined by the SELECT function, and adds it to the sample S. Some or all the neighbors of row r are visited, as determined by the PARENTS and CHILDREN functions, and added to the frontier F. This loop continues until the frontier is empty or some stopping condition, determined by the DONE function, is met.

Algorithm 2 shows an instantiation of Algorithm 1 for RT [29]. As before, ignore the blue lines for now. RT's sampler fills a context window C with tokens sampled via graph traversal. It adds all the parents of a row r and samples w children. When faced with a choice among alternative rows, the sampler chooses at random, without regard to a row's relevance to the prediction task (lines 10 and 18). This type of uninformed random sampling is performed by *all* RFM samplers of which we are aware.

The place where STEER can intervene to fix this problem is the step in which the sampler decides which neighbors of a visited row are added to the frontier (lines 5 and 7 of Algorithm 1). STEER intervenes after these steps to guide the traversal towards rows that are more informative for the prediction task. As described in Section 2, neighbors are reached through two types of links, F→P and P→F, which differ in their fan-out: an F→P traversal reaches at most one parent, whereas a P→F may reach many child rows.

Consider again the Formula 1 database from Section 1 and the task of predicting whether a driver will fail to finish an upcoming race. The seed row in the task table reaches the corresponding row in the `driver` table through its foreign key. That `driver` row, in turn, connects through P→F links to several child tables, including `result`, `standing`, and `qualifying`. As discussed earlier, `result` rows are more directly informative for the task than `standing` or `qualifying` rows. STEER can guide the traversal to favor P→F links leading to `result` over those leading to the other two tables.

The same reasoning applies to F→P links. From a `result` row, for example, following the link to its parent `constructor` row may be informative because some constructors field more reliable cars than others. By contrast, following the link to the parent `race` row is likely less informative, since attributes such as race date or round have little direct bearing on mechanical failure. STEER can therefore favor the F→P link to `constructor` over the one to `race`.

STEER controls graph traversal by applying a *traversal filter* after the sampler identifies the neighbor rows that it wants to add to the frontier. The filter has two variants, one for each traversal direction: an F→P *edge filter* and a P→F *edge filter*. Each filter is a function that is called after the sampler chooses its neighbors. This function decides, for every one of these neighbors, whether the edge from the current row r to that neighbor should be traversed. This decision is made based on the estimated usefulness of the edge to the prediction task. If an edge is filtered out, its destination row is never added to the frontier and therefore never visited. Consequently, none of the neighbors reachable through that row are explored either, so filtering a single edge can prune the entire branch of the traversal rooted at that edge. The blue lines in Algorithms 1 and 2 show the definition of STEER's traversal filter and how it is added to the sampling loop of an RFM sampler. Adding STEER to the sampler of any other RFM will follow the same process.

We modified the samplers of three RFMs—RT, RT-J, and Griffin—to incorporate STEER's traversal filter. An important observation from this implementation effort is that all three models perform sampling within their machine-learning Python code, even though the data being sampled is relational and would ordinarily reside in a SQL DBMS. Adapting three distinct Python-based samplers required substantial model-specific implementation effort. As a

direction for future work, we envision a unified RFM sampler implemented using SQL and exposed to different RFMs through a common API. Such an approach would simplify sampler implementation and make sampling strategies easier to modify and extend. It would also allow RFMs to operate directly on databases managed by a DBMS, without first exporting the data into a separate machine-learning pipeline, while benefiting from the query processing and optimization capabilities of the DBMS [16].

We now turn to the SteerFilter function. A key aspect of this function is that it makes the decision to retain or discard an edge in a *probabilistic* manner. A function $\Phi_{T,d}(e)$ specifies, for prediction task T and edge e in direction d (F→P or P→F), the probability that e is relevant to T. SteerFilter flips a coin with this probability (i.e., conducts a Bernoulli trial) to decide whether to keep the edge e. Thus, edges assigned higher relevance are more likely to be retained, while less relevant edges are more likely to be pruned. Since $\Phi_{T,d}(e)$ is the probability that edge e is included in the sample, we term it the *inclusion probability*.

We made the design decision to drop edges in STEER probabilistically based on the following observation: *edge relevance is rarely binary*. Although one could develop a hard classification of edges into either informative or uninformative for a given task, in practice edges are *informative to varying degrees*. Even edges that are not directly relevant to the prediction target may still contribute useful signal. Returning to the Formula 1 example, `standing` and `qualifying` rows may provide some information for predicting whether a driver will fail to finish, even if they are less informative than `result` rows. Rather than eliminating all `standing` and `qualifying` rows, it is therefore preferable to reduce how often rows from these tables are sampled while preserving the more informative `result` rows. This finer-grained control is naturally achieved through probabilistic edge filtering.

The function $\Phi_{T,d}(e)$ is task-specific, so STEER must define two such relevance functions for every prediction task T, one for each of the F→P and P→F directions. It is clear that the effectiveness of STEER depends critically on defining good $\Phi_{T,d}(e)$ functions that accurately reflect the semantic relevance of an edge to a prediction task. We turn to this important topic next.

4 DETERMINING RELEVANCE

4.1 Estimating Semantic Relevance with LLMs

In this section, we discuss how to decide if an edge is relevant to a given prediction task. In Section 5, we turn this decision into a $\Phi_{T,d}(e)$ function. Deciding whether an edge is worth traversing is fundamentally a question of meaning: does the table at the other end of this edge describe something related to the prediction target? This kind of question can often be answered without inspecting a single row of the database. Table and column names are typically chosen to describe the real-world concepts they represent, and that description alone is often enough to judge whether a concept is likely to matter for a given prediction. Semantic relevance of this kind has been used in other settings such as training task-specific RDL models [6].

LLMs are well suited to producing this kind of semantic judgment. They have been shown to interpret schema-level information effectively in other data management tasks that similarly require judging correspondences or relevance from names and descriptions alone, such as schema matching [24] and text-to-SQL [21]. LLMs have also been used to directly guide sampling for inference in GNNs [2]. We use an LLM for this purpose in STEER, through a purpose-built prompting procedure that elicits a structured relevance judgment for every traversal edge in the schema.

The LLM in STEER is given the full database schema, and it judges *schema edges*. Each foreign-key column defines one schema edge e, which the sampler can traverse in two directions $d \in \{\text{F}\rightarrow\text{P}, \text{P}\rightarrow\text{F}\}$. A table with several foreign keys to the same parent defines several edges, and a foreign key that references its own table defines a self-referential edge. We rank edges rather than tables because the value of reaching a table depends on where the sampler comes from: two edges that lead to the same table can differ in relevance depending on the table from which each is reached.

4.2 Prompting Strategy

A natural way to obtain a judgment about the relevance of an edge is to use a simple prompt: give the LLM the database schema and the prediction task, ask it directly which edges are relevant, and accept a binary include-or-exclude answer for each candidate. Our experiments show that modern LLMs give reasonable results even with this simple prompt (Section 8.4). However, there is room for improvement, and the prompt used by STEER builds on this simple prompt in two ways.

First, it asks the LLM to reason in explicit steps, first about the database schema, then about the prediction task, and only then about how each edge relates to them, rather than answering directly from the raw schema and task description. Judging the relevance of an edge depends on understanding the schema as a whole and the task, so we expect this ordering to help. Guiding a language model through intermediate reasoning steps has been shown to improve the reliability of its outputs on tasks that require combining several pieces of information before reaching a conclusion [34].

Second, instead of a binary decision, the prompt asks the LLM to assign each edge to one of four ordered *tiers* of relevance: *Very Strong*, *Strong*, *Moderate*, or *Weak*. A binary judgment cannot distinguish an edge that is worth following occasionally from one that is not worth following at all from one that is worth following always. A tiered ranking can make such a fine-grained distinction. Comparative, graded judgments of this kind have been found to yield more consistent outputs from LLMs than isolated binary ones [26].

The STEER prompt is organized into four components: (1) a task objective instructing the LLM to rank edges by their semantic relevance to the prediction target, (2) the reasoning steps described above, (3) a set of rules constraining the output to valid schema elements, to prevent the LLM from hallucinating tables that do not exist, and (4) the concrete database and task definition, including table names, column names, primary and foreign keys, and the target column. The LLM is consulted once for each new database and prediction task. Two prompts are issued, one for each traversal direction, and the resulting configuration is reused across every subsequent prediction on that task, so the cost of querying the LLM is incurred once per task rather than once per prediction. The simple and STEER prompt templates are listed in Appendix A.

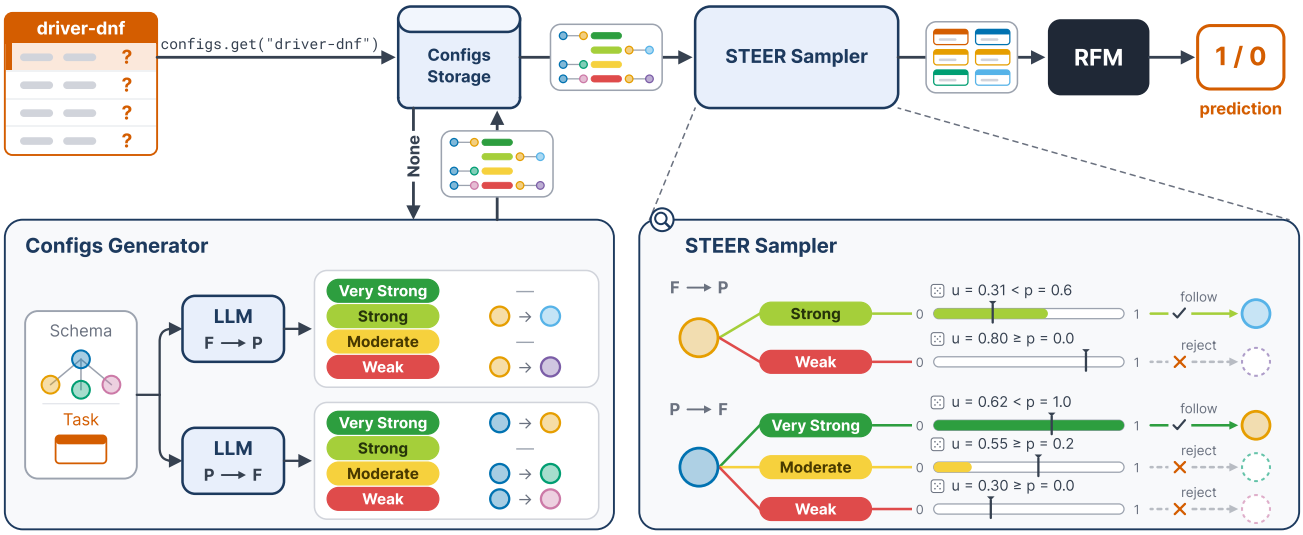


Figure 2: The STEER pipeline. A prediction request is a task table whose rows are prediction targets. STEER looks up the configuration of the task in the Configs Storage. If the storage returns `None`, the Configs Generator generates a configuration with two LLM calls, one per traversal direction, that rank the schema edges into four relevance tiers. STEER saves this configuration in the Configs Storage, so that future invocations of this task will not result in LLM calls. The STEER Sampler then produces a sample using filtered traversal based on the configuration (Algorithm 1). The resulting sample is given to the RFM, which uses it for inference and outputs a prediction.

5 INCLUSION PROBABILITIES

The LLM produces a configuration that assigns a relevance tier to every schema edge in each traversal direction. What remains is converting these tiers into inclusion probabilities $\Phi_{T,d}(e)$ for the traversal filters (Section 3). We adopt the following simple approach for this conversion: for every task, we assign an inclusion probability to each relevance tier in a given traversal direction. Since we have four relevance tiers and two traversal directions, this is a total of eight probabilities for each task. All edges in a given tier get assigned that tier's inclusion probability.

Thus, we define two functions: (1) $\text{tier}_{T,d}(e)$, which is the relevance tier that the LLM assigns to edge e in direction d for task T, and (2) $\pi(v)$ which maps relevance tier v to a probability. Using these two functions, we define $\Phi_{T,d}(e) = \pi\big(\text{tier}_{T,d}(e)\big)$.

We also refer to the tier-to-probability mapping function $\pi(v)$ as the *probability profile* of a task. It is possible to tune the probability profile for every task, but this is costly since it requires performing a search over the possible values of the mapping function, and running the full RFM inference pipeline for each point in this search. Instead, our default approach in STEER is to use a *fixed* probability mapping for all tasks and navigation directions. Specifically, *Very Strong* maps to 1.0, *Strong* to 0.6, *Moderate* to 0.2, and *Weak* to 0.0. We show in Section 8.6 that this choice yields good results, and probability tuning is warranted in only limited cases.

We can now present another reason for adopting probabilistic edge filtering, in addition to the edge informativeness reason discussed in Section 3: probabilistic filtering makes STEER more robust to errors made by the LLM. A binary approach, where all edges in high relevance tiers are sampled and all edges in low relevance tiers are dropped, increases the impact of LLM errors. An edge placed one tier too low can be permanently cut off from the sample, however useful it might actually have been, while an edge placed one tier too high would permanently take up room in the context that it does not deserve. The probabilistic approach mitigates this effect and turns the LLM's ranking from a fixed verdict into a graded influence on traversal: the higher the tier of an edge, the larger the share of its candidates that the sampler follows, and an edge assigned to a neighboring tier by mistake is followed somewhat more or less often, rather than being included or excluded outright.

6 STEER PIPELINE

This section puts the components of STEER together into an end-to-end pipeline, shown in Figure 2. The pipeline starts when an RFM using STEER is presented with a prediction request in the form of a task table whose rows are prediction targets. The pipeline has two parts that run at different frequencies: obtaining the configuration using the LLM, which happens once per database and task, and sampling with the configuration, which happens for every row of the task every time the task is invoked.

STEER conserves calls to the LLM by obtaining a configuration for each task and database once, and reusing the configuration on subsequent invocations of the same task. To accomplish this, STEER stores all computed configurations in a *Configs Storage*, which is a key-value store where the key is a task and the value is the STEER configuration for that task. For each prediction request, STEER first performs a lookup in the Configs Storage using the

task table as the key (for example, `configs.get("driver-dnf")`). If a configuration is found for the task, STEER passes it directly to the sampler. If the storage returns `None`, this indicates the first time a task is seen, and the *Configs Generator* creates the configuration. It issues two LLM calls, one for each traversal direction (F→P and P→F), each with the schema and the task definition (Section 4.2), and creates a configuration consisting of a relevance tier for every schema edge. The configuration is saved in the Configs Storage before it is passed on to the sampler. A configuration depends only on the schema and does not change for a given task, regardless of the rows being predicted. The STEER sampler (Algorithm 1) then uses the configuration and probability profile to build a sample for each row in the task table, and passes this sample to the RFM which produces a prediction. The cost of STEER is, therefore, a lookup in the Configs Storage for a previously seen task, and two LLM calls when a task is seen for the first time.

7 EXTENDING STEER: COLUMN FILTER

The mechanism that STEER uses to filter edges applies equally well to a second point in sampling: once a row has been reached, not every one of its columns is equally useful to include in the sample. A `result` row, for instance, contains a column recording the outcome of a driver's race, including whether they failed to finish, which is very relevant to a did-not-finish prediction. It also contains columns such as a driver's starting grid position or lap time, which are useful but comparatively less informative for this particular task. The row also contains columns such as the result identifier or a comment that carry no predictive signal at all. Writing every column of a visited row to the sample without distinction spends context budget on cells that do little to help the prediction. Unlike the edge filters, which reduce the sample by changing the rows that are visited, the column filter changes what is written to the context once a row has already been visited: withholding a column removes the context cells that column would have contributed, which shrinks the context and leaves room for more informative columns, or for more rows, within the same context length.

The column filter applies only to transformer-based RFMs such as RT and RT-J, in which each cell of a row is a separate element of the context, so a single column can be individually withheld. In graph-based RFMs such as Griffin, a row is represented as a single node, so there is no per-column unit to remove.

We address columns the same way we address traversal. An LLM ranks every column of every table into one of the same four tiers, using the prompting strategy described in Section 4.2, and the tier of a column c is mapped to an inclusion probability $\Phi_{T,d}(c) = \pi\big(\text{tier}_{T,d}(c)\big)$, using the same mapping π as for edges. When a row is visited, each of its columns is written to the context with the probability defined by $\Phi_{T,d}(c)$, and the writing decision for each column is decided independently by a coin flip (Bernoulli trial).

We identify a relatively common case where the column filter can fall short: the LLM derives its ranking entirely from the column name and does not inspect the values that the column actually stores, so it may recommend a column where a large fraction of the values are actually nulls. Such a column would consume context budget while offering little practical signal, a fact that is not knowable from the column name. This phenomenon reflects a more general pattern in how LLM judgments are best used: they are most reliable when used to guide a downstream procedure grounded in the actual data rather than being trusted directly [23].

We correct for this with a lightweight statistical pass over the database: for every column c, we compute its *null rate* from the data, defined as the fraction of rows in the table for which this column is null. If the null rate of a column exceeds a fixed threshold, we demote it to the *Moderate* tier, regardless of the tier the LLM originally assigned it. We settle on *Moderate* because a column that is mostly null may still carry some signal in the rows where a value happens to be present, so its relevance is better described as uncertain rather than clearly high or clearly low. Like other aspects of the relevance configuration, this correction is computed once per new task and subsequently reused.

The column filter is independent of the edge filters and can be used with or without them. Since the column filter works only for transformer-based RFMs and not graph-based RFMs, and since STEER focuses primarily on the navigation during sampling, we consider the column filter to be an *optional extension* to STEER. We show in Section 8.5 that this extension can be useful, but it is not as robust as the traversal element of STEER.

8 EXPERIMENTS

8.1 Experimental Setup

Hardware. All experiments were conducted on a GPU server with 8 NVIDIA H100 GPUs with 80 GB of memory each, 2 AMD Genoa CPUs, and 2.3 TB of RAM.

Benchmark. We evaluate on the RelBench benchmark [31], using every database from RelBench v1 except `rel-event`, which we exclude due to reported temporal leakage issues [29], together with two databases added in RelBench v2 [10], `rel-arxiv` and `rel-ratebeer`. This gives a total of 24 tasks across 8 databases, spanning both classification and regression.

RFM Details. We evaluate STEER on three relational foundation models: RT [29], RT-J [30], and Griffin [33]. We evaluate on RT since the best RFM accuracy at the time of writing is achieved by models from the RT family. RT-J is the most recent model from this family, and the most accurate in recent evaluations. Griffin is an older, less accurate model, but we include it to demonstrate that STEER works with any RFM, including GNN-based models like Griffin, not just transformer-based models like RT or RT-J.

For RT and RT-J, we used the code and model checkpoints provided by the authors.[1] The provided RT checkpoints do not include the new `rel-arxiv` and `rel-ratebeer` databases, so we trained RT from scratch using the code provided by the authors on a corpus that includes these two databases in addition to the six databases on which RT was originally trained. We use the *leave-one-DB-out* curriculum from the RT paper [29]: when conducting inference on a database, we use a model checkpoint trained on the remaining 7 databases, so the inference database is not seen by the model during training and the predictions are zero-shot. RT-J is trained on a very large corpus, and we use the model checkpoint provided by the authors. For Griffin, we use a model checkpoint provided by the RT authors via personal communication, since they modified the model to enable zero-shot predictions.

[1] https://github.com/stanford-star/relational-transformer

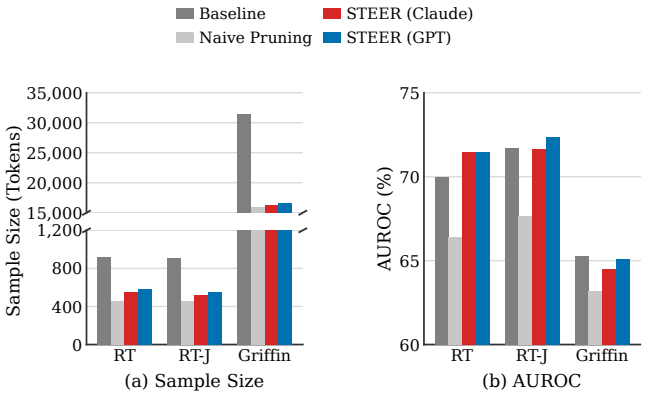


Figure 3: Classification results: average sample size and test accuracy (AUROC).

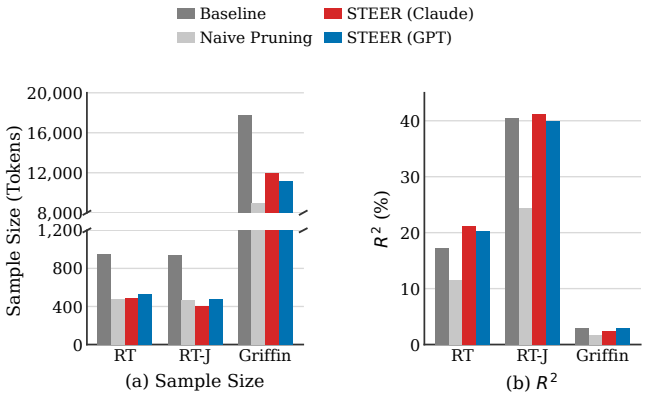


Figure 4: Regression results: average sample size and test accuracy (R^2).

We modified the samplers of all three RFMs to use STEER during inference. RT and RT-J are both configured to sample with a context length of 1024 tokens. Griffin is configured to sample with a fanout of 20 and a traversal depth of 2 hops. All our code and model checkpoints are available.[2]

LLMs. We use two frontier LLMs to produce STEER's tiered ranking: Anthropic's Claude Sonnet 5 and OpenAI's GPT o3. Unless otherwise noted, the LLMs are queried through their APIs using the smart prompting strategy described in Section 4.2, at high reasoning effort. STEER requires two LLM API calls the first time a task is seen, one for F→P and one for P→F.

Accuracy Metrics. Classification tasks are evaluated with AUROC and regression tasks with R^2. Both are reported as percentages (i.e., multiplied by 100), with higher values indicating better accuracy for both metrics. Every reported result, including sample size, is an average over an evaluation budget of 1280 items per task. For all models, one inference run on this number of items typically takes less than 5 minutes on 1 GPU on our server.

Sampler Configurations and Sample Size. We compare four sampler configurations: *Baseline* constructs a sample through unmodified random sampling at each model's default budget. *Naive pruning* constructs a sample the same way, then discards half of the resulting rows uniformly at random, isolating the effect of simply shrinking the sample from the effect of carefully choosing rows to keep. *STEER (Claude)* and *STEER (GPT)* apply STEER's traversal filter, using tiered rankings from Claude and GPT, respectively.

For RT and RT-J, we report sample size as the number of tokens placed in the context window by the sampler. Griffin does not tokenize its input the same way, so we instead report an estimated token count, obtained by converting Griffin's sampled node count using the average number of columns per row across the database.

Default Probability Profile. Unless otherwise noted, STEER uses a single, universal tier-to-probability mapping applied identically across every task, model, and navigation direction: *Very Strong* maps to 1.0, *Strong* to 0.6, *Moderate* to 0.2, and *Weak* to 0.0.

[2]https://github.com/lids-lab/steer

Detailed Results. This section reports the summary of the experimental results. The full task-by-task results are reported in Appendix B.

8.2 Main Result: Effectiveness of STEER

We compare the four sampler configurations described above on all 24 tasks. Figures 3 and 4 report the average sample size for each configuration and the corresponding average test accuracy on classification tasks and regression tasks, respectively.

Let us examine the sampler configurations one by one. The baseline sampler for RT and RT-J does not, on average, fill the entire 1024-token context window, indicating that it often exhausts the available eligible tokens before reaching the context limit. The resulting sample sizes are similar across the two models and across classification and regression tasks, demonstrating the consistency of the sampling procedure. Under baseline sampling, RT-J achieves slightly higher accuracy than RT on classification tasks and much higher accuracy on regression tasks, consistent with RT-J being a larger and more capable model. Griffin, in contrast, uses samples that are roughly an order of magnitude larger than those of RT and RT-J while achieving considerably lower accuracy. This pattern persists across all sampler configurations and is consistent with prior comparisons showing that Griffin underperforms RT [29]. As noted earlier, we include Griffin to demonstrate that STEER applies to GNN-based RFMs as well as transformer-based ones.

Naive pruning reduces sample size by exactly 50% in all cases, by construction, since it always discards half the sample. However, this reduction comes at a substantial cost in prediction accuracy. Naive pruning performs markedly worse than all other sampler configurations. For RT-J, for example, naive pruning reduces classification AUROC by 4.10 percentage points relative to the baseline sampler and reduces regression R^2 by 16.09 percentage points, representing a severe degradation in prediction quality.

STEER also substantially reduces sample size relative to the baseline sampler across all three RFMs and with both LLMs. For RT and RT-J, sample size is reduced by 36.9–43.3% on classification tasks and by 32.8–56.7% on regression tasks. For Griffin, the corresponding reductions are 47.5–48.2% for classification and 32.8–37.0% for regression. These reductions are comparable to the 50% reduction

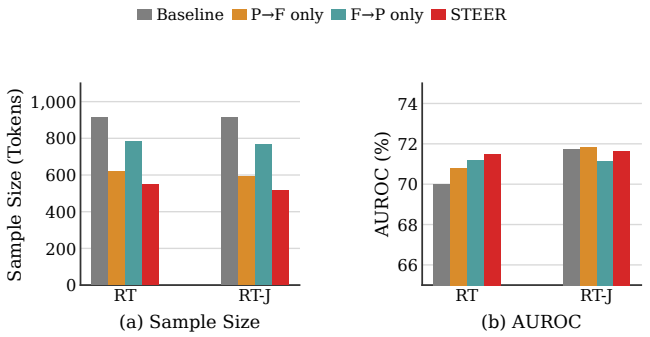


Figure 5: Classification results for direction of traversal ablation: average sample size and test accuracy (AUROC).

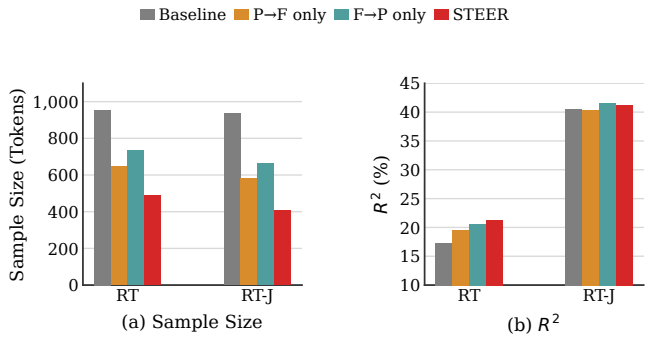


Figure 6: Regression results for direction of traversal ablation: average sample size and test accuracy (R^2).

achieved by naive pruning. In one case, RT-J on regression tasks using Claude, STEER reduces sample size by 56.7%, exceeding the reduction produced by naive pruning.

Crucially, STEER achieves these reductions with little or no loss in prediction accuracy. At worst, STEER incurs only a small decrease relative to the baseline sampler, and in several cases it improves accuracy. For RT, classification AUROC improves by as much as 1.48 percentage points and regression R^2 by as much as 3.89 percentage points. RT-J likewise exhibits accuracy improvements under STEER. We attribute these gains to STEER's ability to remove less useful information from the sample, yielding a more focused and task-relevant inference context. Griffin exhibits only minor accuracy degradation across both task types and both LLM sources, with a maximum decrease of 0.78 percentage points. Its performance therefore remains close to the baseline and substantially better than naive pruning on the same model.

In summary, the central empirical finding of this paper is that **STEER substantially reduces sample size without sacrificing prediction accuracy**. In contrast, naive pruning reduces sample size only at the cost of significant accuracy degradation. These results demonstrate that task relevance is critical when reducing the inference context, and that STEER's LLM-based semantic ranking effectively captures this relevance. We also find that both Claude and GPT are effective sources of semantic information for STEER. Finally, comparing RFMs, we find that RT and RT-J consistently achieve strong prediction accuracy with relatively compact samples, whereas Griffin requires substantially larger samples while achieving lower accuracy.

We next present ablation experiments examining the individual components of STEER. We ran these experiments with all three RFMs and both LLMs and observed the same qualitative patterns throughout. However, because Griffin is substantially weaker than RT and RT-J in our evaluation, and because Claude and GPT produce similar results, we report the remaining experiments only for the transformer-based RT and RT-J models and use Claude as the LLM.

8.3 Ablation: Direction of Traversal

STEER's traversal filter operates on both traversal directions simultaneously: F→P links, which move from a child row to a parent row that it references, and P→F links, which move from a parent row to child rows that reference it (Section 2). This experiment isolates the contribution of each direction by applying STEER's tiered ranking to only one direction at a time while leaving the other unfiltered. The goal is to determine whether one direction contributes more than the other and whether filtering both directions together is more effective than filtering either direction alone.

Figures 5 and 6 report the average sample size and corresponding test accuracy for each traversal direction alone and for STEER, which combines the two directions. The figures show results for classification and regression tasks, respectively.

Filtering either direction alone reduces sample size. P→F filtering reduces the number of tokens by 32.2–35.0%, while F→P filtering yields reductions of 14.5–29.3% across both models and task types. Neither direction alone, however, matches the reduction achieved by filtering both directions together (40.3–56.7%). The two directions are therefore complementary: each targets a different part of the redundancy in a random sample, and combining them removes more than either removes on its own.

Prediction accuracy exhibits a similarly favorable pattern. For RT, filtering P→F links improves accuracy on both classification and regression tasks, while filtering F→P links produces slightly larger improvements. STEER, which filters both directions, yields the largest accuracy gains. For RT-J, all three configurations remain close to baseline accuracy, with only small increases or decreases.

Overall, both traversal directions contribute meaningfully to reducing sample size, and neither direction results in a loss of accuracy. Using both directions together yields the largest reduction in sample size and the highest prediction accuracy.

8.4 Ablation: Prompting Strategy

STEER's smart prompting strategy (Section 4.2) combines structured reasoning steps with a tiered output format. This experiment compares it against simple prompting, which asks for the same relevance judgment directly with no reasoning steps and a binary output. Our goal is to compare the two prompting strategies across multiple LLM variants. Therefore, we conduct this experiment with both Claude and GPT. To investigate whether a powerful LLM will be substantially better than a simpler LLM, we use two settings for the reasoning effort in both models: *High*, the setting used throughout the rest of this paper, and *Default*, the LLM's standard setting. Table 1 reports average test accuracy.

Table 1: Prompting strategy ablation: average test accuracy for classification and regression, comparing simple prompting with STEER's smart prompting. Results are shown for Claude and GPT with high and default reasoning effort. Within each simple/STEER pair, bold indicates the higher accuracy.

RFM	Base.	Claude High		Claude Default		GPT High		GPT Default	
		Simple	STEER	Simple	STEER	Simple	STEER	Simple	STEER
Classification (AUROC, %)									
RT	69.98	71.21	**71.46**	70.58	**71.30**	69.64	**71.44**	70.04	**71.27**
RT-J	71.72	70.83	**71.64**	71.34	**71.52**	70.60	**72.37**	70.77	**71.92**
Regression (R^2, %)									
RT	17.33	17.59	**21.21**	18.05	**21.51**	16.35	**20.31**	17.14	**20.60**
RT-J	40.43	34.12	**41.22**	34.22	**41.03**	33.85	**39.91**	35.61	**39.28**

Table 2: Column filter ablation: average test accuracy and sample size for classification and regression, comparing baseline sampling, STEER without column filter, and STEER with column filter. Bold indicates the highest test score in a row and the highlighted cell indicates the lowest token count.

RFM	Baseline		STEER		STEER + Col. Filter	
	Test ↑	Tok. ↓	Test ↑	Tok. ↓	Test ↑	Tok. ↓
Classification (AUROC, %)						
RT	69.98	915.6	**71.46**	546.6	70.82	467.4
RT-J	71.72	912.6	71.64	517.7	**72.20**	450.6
Regression (R^2, %)						
RT	17.33	954.5	21.21	489.1	**26.60**	409.2
RT-J	40.43	937.2	**41.22**	405.8	38.11	353.9

The main takeaway from the table is that STEER's smart prompting outperforms simple prompting in every setting, thus justifying STEER's use of smart prompting. The gap is modest on classification, typically under one percentage point, indicating that smart prompting has low benefit for these tasks. However, the gap is substantial on regression: RT-J's R^2 under simple prompting falls around 6 percentage points below the baseline across all LLMs, while STEER stays near or above the baseline, thereby providing strong support for using smart prompting.

A second takeaway from the table is that STEER works well with all LLMs. Claude is slightly more accurate than GPT for regression, but both LLMs are sufficiently accurate. There is no big difference between high and default reasoning effort, indicating that STEER is effective even with less powerful, less expensive LLMs.

8.5 Ablation: Column Filter

Section 7 proposes extending STEER's tiered ranking from edges to individual columns, filtering which columns of a visited row are written into the sample. This experiment evaluates the column filter applied on top of STEER's traversal filter. (In this experiment and the next, we return to using only Claude-based rankings.) Table 2 reports average test accuracy and sample size for baseline sampling, STEER without the column filter, and STEER with the column filter. The table reports results for both classification and regression.

As expected, adding the column filter on top of STEER's traversal filter reduces sample size in every case, by an additional 5.6–9.8 percentage points compared to traversal filtering alone. The effect on prediction accuracy is not as uniformly positive across models. On RT, the column filter improves R^2 substantially (5.39 percentage points compared to STEER without the column filter) and leaves classification accuracy close to STEER without the column filter. On RT-J, the column filter improves classification slightly (+0.56 percentage points compared to STEER without the column filter) but drops regression accuracy by 3.11 percentage points.

Overall, the column filter has a directionally positive effect: it improves the sample size in all cases and substantially improves prediction accuracy in some cases. However, there are cases where it hurts prediction accuracy. Therefore, we leave the column filter as a promising extension to STEER that is useful but risky, and thus warrants further investigation as part of future work.

8.6 Ablation: Probability Tuning

Throughout this paper, we have used a single, default tier-to-probability mapping (which we also refer to as a probability profile). This experiment tests whether tuning this mapping on a per-task basis, using a coordinate-ascent search over the validation split, improves upon this default. Because the tuned profile is chosen independently for every task, we report full per-task results in Tables 3 and 4 rather than averages alone.

Probability tuning finds a substantially different probability profile for nearly every task. Tables 3 and 4 show that the benefit of tuning varies widely. Tuning mostly reduces sample size, and it improves accuracy dramatically for some tasks: on rel-ratebeer/user-rating-churn, RT's AUROC rises from 73.02 with the default profile to 82.75 with tuning, and on rel-amazon/item-ltv, its R^2 rises from 4.26 to 36.05.

On the other hand, the accuracy decreases with tuning for some tasks, and RT-J's classification tasks show this clearly: AUROC of rel-trial/study-outcome drops from 52.68 with the default profile to 47.11 once tuned, and rel-stack/user-badge drops from 87.01 to 86.26. This is consistent with a known risk of any validation-based search: it can find a profile that fits the validation split without generalizing well to the test split.

On average, probability tuning improves accuracy for three of the four model/task-type combinations. The only exception is RT-J on classification tasks, due to the two tasks mentioned above for which tuning reduces accuracy. Overall, probability tuning can yield meaningful improvements over the default probability profile, but these gains are task- and model-dependent and are not guaranteed. Moreover, tuning incurs additional cost because it requires

Table 3: Probability tuning ablation: classification tasks (AUROC %, higher is better). For each task and model, we report test accuracy and number of tokens with baseline sampling, STEER with the default probability profile, and STEER with a probability profile tuned for the specific task. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 14 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		STEER		STEER + Prob. Tuning	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-churn	RT	68.20	1023.2	**71.06**	230.2	71.06	230.2
		RT-J	72.26	1018.9	**74.21**	118.7	74.21	118.7
rel-amazon	user-churn	RT	64.14	1024.0	**65.29**	82.5	64.55	71.2
		RT-J	65.45	1021.3	**65.53**	68.1	65.13	58.6
rel-arxiv	paper-citation	RT	79.85	989.6	**79.99**	863.2	79.75	919.2
		RT-J	**80.54**	997.3	79.99	862.7	80.33	922.3
rel-avito	user-clicks	RT	51.92	613.6	63.80	487.5	**66.07**	149.5
		RT-J	**49.64**	612.2	40.52	463.6	41.57	132.6
rel-avito	user-visits	RT	55.98	706.7	**59.47**	374.8	59.47	336.1
		RT-J	63.77	731.7	**64.94**	150.7	63.75	119.1
rel-f1	driver-dnf	RT	78.55	1005.2	**81.40**	999.4	81.38	989.2
		RT-J	**78.50**	1001.5	77.58	995.3	78.36	979.3
rel-f1	driver-top3	RT	89.26	1000.2	**90.70**	1000.2	90.46	775.7
		RT-J	**89.23**	996.6	89.15	996.8	88.96	337.6
rel-hm	user-churn	RT	63.80	1024.0	**65.40**	283.3	65.26	130.2
		RT-J	**63.63**	1020.3	61.49	197.6	62.45	994.2
rel-ratebeer	beer-rating-churn	RT	68.25	1024.0	70.93	1005.8	**71.38**	897.2
		RT-J	66.98	1011.7	74.82	976.7	**75.81**	875.5
rel-ratebeer	brewer-dormant	RT	**71.79**	1024.0	69.54	621.4	70.59	613.0
		RT-J	61.66	1007.8	**61.83**	631.8	61.41	496.5
rel-ratebeer	user-rating-churn	RT	73.38	1024.0	73.02	887.3	**82.75**	361.4
		RT-J	84.29	1013.1	85.72	898.2	**87.37**	1005.0
rel-stack	user-badge	RT	**81.69**	355.6	75.89	134.2	73.67	39.5
		RT-J	**89.53**	357.8	87.01	125.0	86.26	225.8
rel-stack	user-engagement	RT	82.83	981.5	85.15	656.8	**88.64**	115.0
		RT-J	**87.59**	975.0	87.52	737.3	85.54	847.1
rel-trial	study-outcome	RT	50.05	1022.8	48.83	25.7	**52.76**	28.1
		RT-J	51.08	1011.0	**52.68**	25.2	47.11	27.3
Mean	(14 tasks)	RT	69.98	915.6	71.46	546.6	**72.70**	404.0
		RT-J	**71.72**	912.6	71.64	517.7	71.30	510.0
Δ	vs. baseline	RT	—	—	+1.48	-40.3%	+2.72	-55.9%
		RT-J	—	—	-0.08	-43.3%	-0.42	-44.1%

searching over multiple probability profiles. Thus, while probability tuning is worthwhile when the potential gains justify this overhead, the main conclusion from this experiment is that the default probability profile provides a robust, effective, and inexpensive choice that performs well across models and tasks.

9 RELATED WORK

Relational Deep Learning. Fey et al. [8] introduced an end-to-end framework for predictive modeling on relational databases, establishing relational deep learning (RDL) as a subfield of machine learning. At its core, RDL represents a database as a relational entity graph where each row is a node and every primary–foreign key relationship is an edge. Initial approaches applied heterogeneous graph neural networks directly to these graphs, and more recently advanced message-passing architectures have been proposed to improve expressiveness on relational data [5]. Transformer-based architectures have also been explored as a more powerful alternative to GNN message passing, achieving stronger performance on relational benchmarks [7, 25]. A fundamental limitation shared by all these approaches is that they are schema-specific: the model is

Table 4: Probability tuning ablation: regression tasks (R^2 %, higher is better). For each task and model, we report test accuracy and number of tokens with baseline sampling, STEER with the default probability profile, and STEER with a probability profile tuned for the specific task. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 10 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		STEER		STEER + Prob. Tuning	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-ltv	RT	4.59	957.2	4.26	203.5	**36.05**	174.3
		RT-J	59.31	951.5	**60.17**	108.4	60.17	108.4
rel-amazon	user-ltv	RT	15.20	1023.2	**26.26**	100.8	25.73	70.3
		RT-J	**61.52**	1021.3	57.19	86.9	54.41	58.6
rel-arxiv	author-publication	RT	22.77	740.5	**24.75**	491.4	22.57	46.9
		RT-J	18.46	728.8	**33.45**	474.4	33.45	474.4
rel-avito	ad-ctr	RT	4.19	1024.0	5.88	999.4	**9.04**	675.2
		RT-J	2.44	1020.8	-3.02	994.6	**2.63**	149.4
rel-f1	driver-position	RT	45.31	982.6	44.69	982.6	**46.01**	982.6
		RT-J	55.40	978.9	55.34	978.4	**55.45**	975.6
rel-hm	item-sales	RT	21.19	751.5	**23.74**	313.4	23.74	313.4
		RT-J	44.12	729.3	**44.30**	186.6	44.30	186.6
rel-ratebeer	user-rating-count	RT	12.91	1024.0	19.47	420.2	**21.74**	350.9
		RT-J	40.83	1013.0	45.06	272.6	**46.15**	178.3
rel-stack	post-votes	RT	28.75	1003.0	**29.18**	413.5	29.18	942.6
		RT-J	**86.08**	989.1	85.25	139.5	84.73	512.1
rel-trial	site-success	RT	14.78	1015.3	34.83	944.9	**34.96**	935.0
		RT-J	43.54	924.8	**46.66**	794.8	46.18	789.8
rel-trial	study-adverse	RT	**3.57**	1024.0	-0.92	21.7	1.38	410.8
		RT-J	-7.37	1014.1	-12.19	22.0	**-1.07**	581.7
Mean	(10 tasks)	RT	17.33	954.5	21.21	489.1	**25.04**	490.2
		RT-J	40.43	937.2	41.22	405.8	**42.64**	401.5
Δ	vs. baseline	RT	—	—	+3.89	-48.8%	+7.71	-48.6%
		RT-J	—	—	+0.79	-56.7%	+2.21	-57.2%

tightly coupled to a particular database schema and cannot transfer to new, unseen databases.

Relational Foundation Models. To overcome the schema-specificity of RDL, recent work has begun developing relational foundation models capable of zero-shot generalization across databases. KumoRFM and its successor KumoRFM-2 [9, 15] propose a graph-transformer-based relational foundation model with both zero-shot and fine-tuning capabilities, but their implementation details and pretraining procedure are not publicly available. A separate line of work applies single-table tabular foundation models [12, 13, 27, 28] to relational data by first flattening the database into a single feature table; RDBLearn [38] follows this strategy, aggregating relational neighborhoods into engineered features before applying an in-context tabular learner, an effective approach that nonetheless discards the relational structure a native model can exploit directly. Two architecturally distinct open-source relational foundation models represent the current state of the art. Griffin [33] pretrains a graph-centric architecture on a mixture of tabular and relational datasets, combining table-level encoders with graph neural network propagation across tables. The Relational Transformer (RT) [29] instead represents every database cell as an individual token, introducing cell-level tokenization, task table prompting, and a novel Relational Attention mechanism to enable zero-shot inference across heterogeneous schemas. Two recent follow-ups retain this same cell-level architecture while changing how it is pretrained: RT-J [30] introduces a new large-scale pretraining recipe aimed at label-efficient few-shot prediction, and RT-PluRel [18] pretrains on synthetic relational databases generated by the PluRel framework. One of the defining features of STEER is that it can work with any RFM since its decisions are at the level of which foreign-key link to follow during sampling. We demonstrated this flexibility by evaluating STEER on three RFMs with different architectures and different samplers.

Sampling Strategies for Graph-Based Models. A rich body of work has developed neighborhood sampling methods for GNNs,

primarily motivated by training efficiency. Hamilton et al. [11] introduced uniform neighbor sampling as part of GraphSAGE, enabling scalable aggregation over sampled local neighborhoods. Ying et al. [39] extended this idea with random-walk-based importance sampling in PinSage, biasing neighborhood selection toward more relevant nodes. Lao and Cohen [19] similarly use path-constrained random walks to guide relational retrieval along specified relation paths. In the RDL setting, Khan and Aboulnaga [16] explore SQL-based sampling within relational database systems as part of a broader vision for performing RDL computation inside the DBMS. More recently, Afandi et al. [2] proposed using LLMs to generate reusable query templates that extract semantically relevant subgraphs for GNN inference on large knowledge graphs, providing semantic guidance at inference time. Ding et al. [6] introduce a structure-semantic collaborative sampling strategy for a task-specific relational graph transformer that filters irrelevant semantic information while preserving structural connectivity. Our work addresses a complementary setting by introducing a semantically informed traversal filter for inference-time sampling in a zero-shot relational foundation model, where no task-specific training data is available to learn what is informative.

LLMs for Relational Data. Large language models have been applied to relational data in several ways. Wydmuch et al. [36] demonstrate that even a straightforward application of LLMs to relational databases—by serializing the database subgraph into flat text—achieves competitive performance on RelBench tasks, establishing LLMs as a promising baseline. Wu et al. [35] improve on this by introducing Rel-LLM, a hybrid architecture that uses a GNN encoder to generate structured relational prompts for LLMs within a retrieval-augmented generation framework, preserving relational structure rather than serializing to flat text. Beyond end-to-end prediction, LLMs have also been shown to reason effectively about schema semantics from attribute names and descriptions alone, both in the context of schema matching [24] and in text-to-SQL systems that link natural-language questions to the correct tables and columns purely from their names [21]. Our approach differs fundamentally from all of these prior approaches: rather than using LLMs to make predictions, match schemas, or decide how much schema to expose to another model, we use them as a semantic oracle that produces a tiered relevance ranking of a database schema's traversal edges to steer sampling at inference time.

10 CONCLUSION

We presented STEER, a semantically informed sampling approach for reducing RFM inference cost. STEER uses an LLM to assess the relevance of schema edges to a prediction task and converts these judgments into traversal probabilities for sample construction. The resulting relevance configuration is reused across predictions, amortizing the cost of the LLM calls. Across three RFMs, STEER substantially reduces inference context size while preserving, and in several cases improving, prediction accuracy. In contrast, uninformed random pruning causes substantial accuracy degradation, showing that effective context reduction requires reasoning about which parts of the relational structure are relevant to the task.

Several directions remain for future work. First, integrating STEER into multiple RFMs highlights the value of a common SQL-based sampling abstraction that can be shared across models. Second, the optional column filter suggests that semantic guidance can extend beyond edge traversal, although its mixed results indicate that more robust methods are needed, potentially combining schema semantics with lightweight data statistics. Finally, task-specific probability profiles can sometimes outperform the default profile, motivating inexpensive methods for adapting traversal probabilities without costly search. More broadly, applying semantic guidance to other RFM architectures and sampling strategies may yield further reductions in inference cost without sacrificing predictive accuracy.

REFERENCES

[1] Dennis Abts, Jonathan Ross, Jonathan Sparling, Mark Wong-VanHaren, Max Baker, Tom Hawkins, Andrew Bell, John Thompson, Temesghen Kahsai, Garrin Kimmell, Jennifer Hwang, Rebekah Leslie-Hurd, Michael Bye, E. R. Creswick, Matthew Boyd, Mahitha Venigalla, Evan Laforge, Jon Purdy, Purushotham Kamath, Dinesh Maheshwari, Michael Beidler, Geert Rosseel, Omar Ahmad, Gleb Gagarin, Richard Czekalski, Ashay Rane, Sahil Parmar, Jeff Werner, Jim Sproch, Adrian Macias, and Brian Kurtz. 2020. Think Fast: A Tensor Streaming Processor (TSP) for Accelerating Deep Learning Workloads. In *Proceedings of the ACM/IEEE Annual International Symposium on Computer Architecture (ISCA)*. 145–158.

[2] Waleed Afandi, Hussein Abdallah, Ashraf Aboulnaga, and Essam Mansour. 2026. An LLM-Guided Query-Aware Inference System for GNN Models on Large Knowledge Graphs. In *Proceedings of the 42nd IEEE International Conference on Data Engineering (ICDE)*. 1675–1688.

[3] Amazon Web Services. [n.d.]. AWS Inferentia: High-Performance Machine Learning Inference Accelerators. https://aws.amazon.com/ai/machine-learning/inferentia/.

[4] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 785–794.

[5] Tianlang Chen, Charilaos Kanatsoulis, and Jure Leskovec. 2025. RelGNN: Composite Message Passing for Relational Deep Learning. In *Proceedings of the International Conference on Machine Learning (ICML)*. 8296–8312.

[6] Zezhong Ding, Jin Li, Xugang Wang, and Xike Xie. 2026. Gaussian Relational Graph Transformer. *CoRR* abs/2605.15575 (2026).

[7] Vijay Prakash Dwivedi, Sri Jaladi, Yangyi Shen, Federico Lopez, Charilaos Kanatsoulis, Rishi Puri, Matthias Fey, and Jure Leskovec. 2026. Relational Graph Transformer. In *Proceedings of the International Conference on Learning Representations (ICLR)*. 154348–154369.

[8] Matthias Fey, Weihua Hu, Kexin Huang, Jan Eric Lenssen, Rishabh Ranjan, Joshua Robinson, Rex Ying, Jiaxuan You, and Jure Leskovec. 2024. Position: Relational Deep Learning—Graph Representation Learning on Relational Databases. In *Proceedings of the International Conference on Machine Learning (ICML)*. 13592–13607.

[9] Matthias Fey, Vid Kocijan, Federico López, Jan Eric Lenssen, and Jure Leskovec. 2025. *KumoRFM: A Foundation Model for In-Context Learning on Relational Data*. Technical Report. Kumo AI.

[10] Justin Gu, Rishabh Ranjan, Charilaos I. Kanatsoulis, Haiming Tang, Martin Jurkovic, Valter Hudovernik, Mark Znidar, Pranshu Chaturvedi, Parth Shroff, Fengyu Li, and Jure Leskovec. 2026. RelBench v2: A Large-Scale Benchmark and Repository for Relational Data. In *Proceedings of the International Conference on Learning Representations (ICLR) Workshop on Foundation Models for Tabular and Structured Data*.

[11] William L. Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*. 1024–1034.

[12] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. 2023. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*.

[13] Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeister, and Frank Hutter. 2025. Accurate Predictions on Small Data with a Tabular Foundation Model. *Nature* 637, 8045 (2025), 319–326.

[14] Coleman Richard Charles Hooper, Sehoon Kim, Hiva Mohammadzadeh, Monishwaran Maheswaran, Sebastian Zhao, June Paik, Michael W. Mahoney, Kurt Keutzer, and Amir Gholami. 2025. Squeezed Attention: Accelerating Long Context Length LLM Inference. In *Proceedings of the 63rd Annual Meeting of the*

Association for Computational Linguistics (ACL). 32631–32652.

[15] Valter Hudovernik, Federico Lopez, Vid Kocijan, Akihiro Nitta, Jan Eric Lenssen, Jure Leskovec, and Matthias Fey. 2026. KumoRFM-2: Scaling Foundation Models for Relational Learning. *CoRR* abs/2604.12596 (2026).

[16] Fahim Shahriar Khan and Ashraf Aboulnaga. 2025. A Vision for SQL-Based Relational Deep Learning. In *Proceedings of the International Workshop on Tabular Data Analysis (TaDA), co-located with the International Conference on Very Large Data Bases (VLDB)*.

[17] Vid Kocijan, Jinu Sunil, Jan Eric Lenssen, Viman Deb, Xinwei He, Federico Reyes Gomez, Matthias Fey, and Jure Leskovec. 2026. Predictive Query Language: A Domain-Specific Language for Predictive Modeling on Relational Databases. In *Proceedings of the International Workshop on Tabular Data Analysis (TaDA), co-located with the International Conference on Very Large Data Bases (VLDB)*.

[18] Vignesh Kothapalli, Rishabh Ranjan, Valter Hudovernik, Vijay Prakash Dwivedi, Johannes Hoffart, Carlos Guestrin, and Jure Leskovec. 2026. PluRel: Synthetic Data Unlocks Scaling Laws for Relational Foundation Models. In *Proceedings of the Forty-Third International Conference on Machine Learning (ICML)*.

[19] Ni Lao and William W. Cohen. 2010. Relational Retrieval Using a Combination of Path-Constrained Random Walks. *Machine Learning* (2010).

[20] Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast Inference from Transformers via Speculative Decoding. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*.

[21] Haoyang Li, Jing Zhang, Hanbing Liu, Ju Fan, Xiaokang Zhang, Jun Zhu, Renjie Wei, Hongyan Pan, Cuiping Li, and Hong Chen. 2024. CodeS: Towards Building Open-Source Language Models for Text-to-SQL. *Proceedings of the ACM on Management of Data* 2, 3 (2024), 1–28.

[22] Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. 2023. Compressing Context to Enhance Inference Efficiency of Large Language Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 6342–6353.

[23] Stefano Marchesin, Matteo Ceccarello, and Gianmaria Silvello. 2026. LLMs as Stratification Signals for KG Accuracy Evaluation. *Proceedings of the VLDB Endowment* 19, 9 (2026), 2005–2018.

[24] Marcel Parciak, Brecht Vandevoort, Frank Neven, Liesbet M. Peeters, and Stijn Vansummeren. 2024. Schema Matching with Large Language Models: an Experimental Study. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases (VLDB)*.

[25] Jakub Peleska and Gustav Sír. 2025. Tabular Transformers Meet Relational Databases. *ACM Transactions on Intelligent Systems and Technology (TIST)* 16, 5 (2025), 115:1–115:24.

[26] Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. 2024. Large Language Models are Effective Text Rankers with Pairwise Ranking Prompting. In *Findings of the Association for Computational Linguistics: North American Chapter of the Association for Computational Linguistics (NAACL) 2024*. 1504–1518.

[27] Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. 2025. TabICL: A Tabular Foundation Model for In-Context Learning on Large Data. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*. 50817–50847.

[28] Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. 2026. TabICLv2: A Better, Faster, Scalable, and Open Tabular Foundation Model. In *Proceedings of the Forty-Third International Conference on Machine Learning (ICML)*.

[29] Rishabh Ranjan, Valter Hudovernik, Mark Znidar, Charilaos I. Kanatsoulis, Roshan Upendra, Mahmoud Mohammadi, Joe Meyer, Tom Palczewski, Carlos Guestrin, and Jure Leskovec. 2026. Relational Transformer: Toward Zero-Shot Foundation Models for Relational Data. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

[30] Rishabh Ranjan, Vignesh Kothapalli, Harshvardhan Agarwal, Charilaos I. Kanatsoulis, Roshan Reddy Upendra, Tom Palczewski, Carlos Guestrin, and Jure Leskovec. 2026. Large-Scale Pretraining Unlocks Few-Shot Prediction for Relational Data. In *Proceedings of the 2nd Workshop on Foundation Models for Structured Data at the International Conference on Machine Learning (ICML)*.

[31] Joshua Robinson, Rishabh Ranjan, Weihua Hu, Kexin Huang, Jiaqi Han, Alejandro Dobles, Matthias Fey, Jan Eric Lenssen, Yiwen Yuan, Zecheng Zhang, and Jure Leskovec. 2024. RelBench: A Benchmark for Deep Learning on Relational Databases. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*.

[32] Ravid Shwartz-Ziv and Amitai Armon. 2022. Tabular Data: Deep Learning is not all you need. *Information Fusion* 81 (2022), 84–90.

[33] Yanbo Wang, Xiyuan Wang, Quan Gan, Minjie Wang, Qibin Yang, David Wipf, and Muhan Zhang. 2025. Griffin: Towards a Graph-Centric Relational Database Foundation Model. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*.

[34] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*.

[35] Fang Wu, Vijay Prakash Dwivedi, and Jure Leskovec. 2025. Large Language Models are Good Relational Learners. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*. 7835–7854.

[36] Marek Wydmuch, Łukasz Borchmann, and Filip Graliński. 2024. Tackling Prediction Tasks in Relational Databases with LLMs. *CoRR* abs/2411.11829 (2024).

[37] Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*.

[38] Linjie Xu, Yanlin Zhang, Quan Gan, Minjie Wang, and David Wipf. 2026. No Need to Train Your RDB Foundation Model. In *Proceedings of the Forty-Third International Conference on Machine Learning (ICML)*.

[39] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 974–983.

[40] Zixuan Zhou, Xuefei Ning, Ke Hong, Tianyu Fu, Jiaming Xu, Shiyao Li, Yuming Lou, Luning Wang, Zhihang Yuan, Xiuhong Li, Shengen Yan, Guohao Dai, Xiao-Ping Zhang, Yuhan Dong, and Yu Wang. 2024. A Survey on Efficient Inference for Large Language Models. *CoRR* abs/2404.14294 (2024).

A PROMPT TEMPLATES

For reproducibility, we report the full text of the two prompt templates used to produce tiered rankings: the smart prompting strategy and the simple prompting baseline (Section 4.2). When creatnig a configuration, each template is used twice for the two navigation directions, once for P→F and once for F→P. Each template is shown in its P→F form in full; the F→P form is identical except where noted. Every prompt ends with a task-specific block, shown here as a placeholder, containing the task name, dataset, task table, task table columns, and the full database definition.

A.1 Smart Prompting Template

```
You are an expert in relational databases and machine
↪ learning on tabular data.

Your job is to rank traversal edges into tiers for a
↪ probabilistic BFS-like primary-to-foreign-key traversal
↪ policy based on semantic reasoning about the database
↪ schema.
In a relational database, a primary-to-foreign-key
↪ relationship means that a table with a primary key is
↪ referenced by other tables via foreign keys. Those
↪ referencing tables are the child tables. When the
↪ sampler is at a row in any table, it can expand outward
↪ into child tables by following those foreign key
↪ references. Each such expansion is a traversal edge -- a
↪ directed connection from the current parent table to a
↪ specific child table via a foreign key. Your goal is to
↪ rank each traversal edge into tiers based on how useful
↪ it is to follow that specific edge during sampling. The
↪ coin toss happens at the edge level when a node is being
↪ pushed into the frontier, so rejected edges are never
↪ enqueued at all.
Your ranking should be based purely on semantic reasoning:
↪ what each table represents, what its columns describe,
↪ how directly that information relates to the prediction
↪ target and the entity being predicted for, and crucially,
↪ what the source table context is when the edge is
↪ crossed.

Follow these steps in order:

Step 1 - Understand the database structure: Read through all
↪ the tables and their columns. For each table, form a
↪ clear understanding of what real-world entity or concept
↪ that table represents and what kind of information each
↪ column holds. Identify all primary keys and foreign keys
↪ and understand how the tables are connected to each
↪ other.

Step 2 - Understand the task: Identify the entity being
↪ predicted for, the target value being predicted, and
↪ what the prediction task is asking you to estimate or
↪ classify.

Step 3 - Build the edge list: Go through each table in the
↪ database and identify which tables have at least one
↪ foreign key defined. For each foreign key column in each
↪ such table, identify its parent table -- the table whose
↪ primary key is referenced by that foreign key column.
↪ Write out all traversal edges as
↪ source->destination:fk_col triples, where source is the
↪ parent table (where the sampler currently is),
↪ destination is the child table (where the sampler would
↪ traverse to), and fk_col is the exact foreign key column
↪ name defined on the destination table. Note that the
↪ arrow in source->destination represents the direction
↪ the sampler travels, not the direction of the foreign
↪ key. The foreign key is always defined on the
↪ destination table pointing back to the source. Each
↪ foreign key column in the database produces exactly one
↪ edge. If a destination table has multiple foreign keys
↪ pointing to the same source, each FK column produces a
↪ separate edge with a distinct fk_col suffix. Tables that
↪ have no foreign keys produce no edges as destinations.
↪ List every edge explicitly before moving on.

Step 4 - Rank edges by semantic relevance: Go through each
↪ edge in your edge list and reason about how useful it is
↪ to follow that specific edge during sampling. The key
↪ question for each edge is: given that the sampler is
↪ currently at a row in the source table, does crossing
↪ this edge to the destination table provide meaningful
↪ signal toward the prediction target and the entity being
↪ predicted for? Also consider what subset of the
↪ destination table's rows you would actually reach by
↪ following this edge from this specific source. A more
↪ targeted subset that directly characterizes the
↪ prediction entity is more valuable than a broad or noisy
↪ one. Two edges that share the same destination table can
↪ have very different values if their source tables
↪ provide different context -- reason about each edge
↪ independently.

Rank all edges into the following tiers: very_strong, strong,
↪ moderate, weak.

very_strong: following this edge from this specific source
↪ will almost always lead to rows that directly describe
↪ or characterize the prediction entity or its outcomes in
↪ a way that is immediately relevant to the target.
strong: following this edge provides important contextual
↪ signal closely related to the target, even if not
↪ directly causal, given the source context.
moderate: following this edge has some plausible
↪ relationship to the target but the connection requires
↪ more than one step of reasoning or is indirect given the
↪ source context.
weak: following this edge leads to rows that describe
↪ entities or events with no plausible direct or indirect
↪ connection to the prediction target given where the
↪ sampler currently is.

Step 5 - Ensure there is at least one edge in your output. If
↪ none qualify, go back to step 4 and select the single
↪ most relevant edge into very_strong.
```

```
Hard constraints: only use edges from your step-3 list (do
↪  not invent edges); use exact table/column names from the
↪  database definition; each edge appears in exactly one
↪  tier.

Return only a JSON object: {"config": {...tiers...},
↪  "reasoning": {"step_1": "...", ..., "step_5": "..."}}.

[Task name / Dataset / Task table / Task table columns /
↪  Database definition inserted here]
```

The F→P variant of this template is identical except for the following. The opening direction is reframed: *"a foreign-to-primary-key relationship means that a table holds a foreign key that references a primary key in another table… the sampler can expand outward into parent tables."* In Step 3, *source* is the child table (the sampler's current position), *destination* is the parent table, and *fk_col* is defined on the source table rather than the destination; self-referential edges, where a table's foreign key references its own primary key, are explicitly valid and must be included. In Step 4, the guidance on targeted subsets is reframed around the destination as a single parent entity: *"a parent row that directly identifies or characterizes the prediction entity is more valuable than one that only provides generic or loosely related context."* The tier definitions, Step 5, hard constraints, output format, and trailing task-context block are otherwise identical to the P→F version.

A.2 Simple Prompting Template

```
You are an expert in relational databases and machine
↪  learning on tabular data.

Your job is to identify which traversal edges are useful for
↪  a probabilistic BFS-like primary-to-foreign-key
↪  traversal policy, based on the database schema and the
↪  prediction task below.
In a relational database, a primary-to-foreign-key
↪  relationship means that a table with a primary key is
↪  referenced by other tables via foreign keys. Those
↪  referencing tables are the child tables. When the
↪  sampler is at a row in any table, it can expand outward
↪  into child tables by following those foreign key
↪  references. Each such expansion is a traversal edge -- a
↪  directed connection from the current parent table to a
↪  specific child table via a foreign key, written as
↪  source->destination:fk_col, where source is the parent
↪  table (where the sampler currently is) and destination
↪  is the child table (where the sampler would traverse to).

List the traversal edges you consider useful for predicting
↪  the target -- edges that lead to rows providing
↪  meaningful signal about the prediction target and the
↪  entity being predicted for.

Use the exact table names and column names from the database
↪  definition. Do not invent edges.

Return only a valid JSON object in this exact format, with
↪  no explanation, no markdown, and nothing else:
{"edges": ["source_a->destination_a:fk_col_x",
↪  "source_b->destination_b:fk_col_y"]}

[Task name / Dataset / Task table / Task table columns /
↪  Database definition inserted here]
```

The F→P variant reframes the opening direction the same way as the smart prompt's F→P variant: *"a foreign-to-primary-key relationship means that a table holds a foreign key that references a primary key in another table… the sampler can expand outward into parent tables… source is the child table, destination is the parent table."* Self-referential edges are again explicitly valid. The instruction line, hard constraints, JSON output format, and trailing task-context block are otherwise identical to the P→F version.

Every edge returned by a simple prompt is placed entirely into the *very_strong* tier, with the remaining three tiers left empty; the resulting configuration is then evaluated using a profile that zeroes every tier except *very_strong*. This is the mechanical difference between smart and simple prompting: a graduated, four-tier ranking produced through explicit step-by-step reasoning, versus a binary include-or-exclude decision with no reasoning scaffold.

B FULL PER-TASK RESULTS

We report the full per-task results for every experiment in Section 8. All values are averaged using the same evaluation procedure described in Section 8.1.

Table 5: Main results: classification tasks (AUROC %, higher is better). Used in Figure 3. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), naive pruning (full-budget sampling, then half the resulting sample discarded at random), STEER with Claude-based ranking, and STEER with GPT-based ranking. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 14 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		Naive Pruning		STEER (Claude)		STEER (GPT)	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-churn	RT	68.20	1023.2	63.90	511.6	71.06	230.2	**71.07**	236.4
		RT-J	72.26	1018.9	66.58	509.9	74.21	118.7	**74.41**	122.6
		Griffin	69.04	10885.8	65.64	5523.0	**69.10**	8601.3	69.07	9057.0
rel-amazon	user-churn	RT	64.14	1024.0	60.39	512.0	**65.29**	82.5	64.55	71.2
		RT-J	65.45	1021.3	62.09	511.2	**65.53**	68.1	65.05	58.6
		Griffin	62.18	6111.9	60.73	3135.9	61.93	5016.0	**62.71**	3370.5
rel-arxiv	paper-citation	RT	79.85	989.6	73.98	494.8	**79.99**	863.2	79.39	945.0
		RT-J	80.54	997.3	74.26	499.2	79.99	862.7	**80.81**	947.5
		Griffin	**74.49**	5488.9	71.13	2800.4	74.04	3234.5	74.47	3363.1
rel-avito	user-clicks	RT	51.92	613.6	53.45	307.2	**63.80**	487.5	63.02	534.4
		RT-J	**49.64**	612.2	44.99	306.6	40.52	463.6	45.30	521.9
		Griffin	50.44	14308.2	49.31	7238.0	50.41	8676.1	**50.47**	9835.4
rel-avito	user-visits	RT	55.98	706.7	57.14	353.6	59.47	374.8	**59.57**	603.8
		RT-J	63.77	731.7	57.02	366.4	64.94	150.7	**65.07**	580.4
		Griffin	59.94	17505.3	59.86	8836.3	60.06	12793.5	**60.16**	9244.0
rel-f1	driver-dnf	RT	78.55	1005.2	73.16	502.6	**81.40**	999.4	80.62	1000.8
		RT-J	**78.50**	1001.5	74.27	501.3	77.58	995.3	78.44	996.9
		Griffin	51.20	44371.9	51.65	22308.0	48.69	26472.8	**51.84**	23747.1
rel-f1	driver-top3	RT	89.26	1000.2	82.90	500.1	**90.70**	1000.2	89.35	1000.2
		RT-J	**89.23**	996.6	85.50	498.8	89.15	996.8	89.18	996.4
		Griffin	61.75	44002.5	60.42	22117.1	60.10	32050.6	**61.88**	37211.2
rel-hm	user-churn	RT	63.80	1024.0	61.61	512.0	**65.40**	283.3	65.40	283.3
		RT-J	**63.63**	1020.3	59.38	510.7	61.49	197.6	61.49	197.6
		Griffin	**60.38**	23628.7	59.97	12012.1	60.05	18939.4	60.05	18939.4
rel-ratebeer	beer-rating-churn	RT	68.25	1024.0	65.93	512.0	**70.93**	1005.8	70.25	1024.0
		RT-J	66.98	1011.7	64.66	506.4	74.82	976.7	**75.75**	1009.6
		Griffin	**62.78**	46983.0	61.34	23765.8	62.28	15021.2	61.84	18660.7
rel-ratebeer	brewer-dormant	RT	**71.79**	1024.0	67.57	512.0	69.54	621.4	70.01	1024.0
		RT-J	61.66	1007.8	55.40	504.4	61.83	631.8	**63.96**	1010.3
		Griffin	**68.74**	145954.7	64.77	73251.0	67.91	41201.3	68.59	46707.2
rel-ratebeer	user-rating-churn	RT	**73.38**	1024.0	72.08	512.0	73.02	887.3	72.37	887.4
		RT-J	84.29	1013.1	83.17	507.1	85.72	898.2	**86.10**	896.6
		Griffin	84.61	32913.9	84.52	16731.8	84.61	28994.1	**84.63**	28485.4
rel-stack	user-badge	RT	**81.69**	355.6	75.95	178.1	75.89	134.2	78.77	172.7
		RT-J	**89.53**	357.8	85.35	179.4	87.01	125.0	87.50	177.2
		Griffin	80.69	10608.3	78.86	5418.5	**80.73**	7830.8	80.37	7558.6
rel-stack	user-engagement	RT	82.83	981.5	71.42	490.8	85.15	656.8	**85.80**	270.7
		RT-J	**87.59**	975.0	79.24	488.0	87.52	737.3	86.81	209.9
		Griffin	**76.78**	23592.3	65.38	11910.7	75.69	18442.1	74.56	12756.7
rel-trial	study-outcome	RT	50.05	1022.8	**50.43**	511.4	48.83	25.7	49.99	40.6
		RT-J	51.08	1011.0	**54.89**	506.1	52.68	25.2	53.26	34.0
		Griffin	**50.97**	14699.2	50.93	7491.2	47.42	1193.7	50.42	2798.8
Mean	(14 tasks)	RT	69.98	915.6	66.42	457.9	**71.46**	546.6	71.44	578.2
		RT-J	71.72	912.6	67.63	456.8	71.64	517.7	**72.37**	554.3
		Griffin	**65.29**	31503.9	63.18	15895.7	64.50	16319.1	65.08	16552.5
Δ	vs. baseline	RT	—	—	-3.56	-50.0%	+1.48	-40.3%	+1.46	-36.9%
		RT-J	—	—	-4.10	-49.9%	-0.08	-43.3%	+0.64	-39.3%
		Griffin	—	—	-2.11	-49.5%	-0.78	-48.2%	-0.21	-47.5%

Table 6: Main results: regression tasks (R^2 %, higher is better). Used in Figure 4. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), naive pruning (full-budget sampling, then half the resulting sample discarded at random), STEER with Claude-based ranking, and STEER with GPT-based ranking. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 10 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		Naive Pruning		STEER (Claude)		STEER (GPT)	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-ltv	RT	**4.59**	957.2	3.03	478.7	4.26	203.5	4.26	203.5
		RT-J	59.31	951.5	24.29	476.2	**60.17**	108.4	60.17	108.4
		Griffin	**15.70**	10175.6	13.27	5167.8	15.38	8079.4	15.56	7241.0
rel-amazon	user-ltv	RT	15.20	1023.2	3.42	511.6	**26.26**	100.8	25.29	101.0
		RT-J	**61.52**	1021.3	11.26	511.1	57.19	86.9	56.01	86.7
		Griffin	**-0.49**	5906.3	-7.51	3033.2	-5.45	4048.7	-1.10	5377.0
rel-arxiv	author-publication	RT	22.77	740.5	23.00	370.3	**24.75**	491.4	23.99	387.1
		RT-J	18.46	728.8	7.72	364.8	33.45	474.4	**35.36**	356.6
		Griffin	-1.67	2140.3	**1.99**	1126.1	-1.67	2140.3	-1.06	1937.5
rel-avito	ad-ctr	RT	4.19	1024.0	1.91	512.0	5.88	999.4	**7.70**	1024.0
		RT-J	**2.44**	1020.8	1.67	511.0	-3.02	994.6	-2.77	1018.6
		Griffin	2.39	16593.3	0.51	8379.8	**2.96**	9025.7	2.42	6927.3
rel-f1	driver-position	RT	45.31	982.6	35.16	491.3	44.69	982.6	**46.90**	982.6
		RT-J	55.40	978.9	39.16	490.0	55.34	978.4	**55.55**	978.4
		Griffin	-2.74	36402.7	-2.65	18322.9	-2.59	28600.6	**-2.40**	25318.9
rel-hm	item-sales	RT	21.19	751.5	2.27	375.8	**23.74**	313.4	21.62	640.5
		RT-J	44.12	729.3	32.16	365.0	**44.30**	186.6	42.81	622.8
		Griffin	**5.54**	30219.9	5.37	15307.8	5.48	25540.4	5.48	26478.5
rel-ratebeer	user-rating-count	RT	12.91	1024.0	13.67	512.0	**19.47**	420.2	18.21	432.9
		RT-J	40.83	1013.0	41.83	507.0	**45.06**	272.6	44.66	289.0
		Griffin	-5.74	32905.8	**-5.64**	16727.2	-5.73	23975.0	-5.69	19543.1
rel-stack	post-votes	RT	28.75	1003.0	22.37	501.5	**29.18**	413.5	28.22	677.9
		RT-J	**86.08**	989.1	58.84	495.1	85.25	139.5	81.05	590.5
		Griffin	17.75	24390.1	13.09	12309.2	**18.03**	9879.3	17.50	10707.1
rel-trial	site-success	RT	14.78	1015.3	8.52	507.6	**34.83**	944.9	22.89	787.5
		RT-J	43.54	924.8	32.96	463.0	**46.66**	794.8	33.44	699.6
		Griffin	1.29	6860.1	0.40	3575.4	1.29	6860.1	**1.33**	6099.5
rel-trial	study-adverse	RT	3.57	1024.0	1.86	512.0	-0.92	21.7	**3.97**	77.2
		RT-J	-7.37	1014.1	**-6.40**	507.6	-12.19	22.0	-7.18	45.2
		Griffin	-2.52	11558.5	-2.29	5923.7	-2.93	947.0	**-1.92**	1944.4
Mean	(10 tasks)	RT	17.33	954.5	11.52	477.3	**21.21**	489.1	20.31	531.4
		RT-J	40.43	937.2	24.35	469.1	**41.22**	405.8	39.91	479.6
		Griffin	2.95	17715.3	1.65	8987.3	2.48	11909.7	**3.01**	11157.4
Δ	vs. baseline	RT	—	—	-5.80	-50.0%	+3.89	-48.8%	+2.98	-44.3%
		RT-J	—	—	-16.09	-49.9%	+0.79	-56.7%	-0.52	-48.8%
		Griffin	—	—	-1.30	-49.3%	-0.47	-32.8%	+0.06	-37.0%

Table 7: Direction of traversal ablation: classification tasks (AUROC %, higher is better). Used in Figure 5. Ablation results use only the RT and RT-J models and the Claude LLM. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), P→F-only traversal, F→P-only traversal, and the STEER traversal filter that combines P→F and F→P. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 14 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		P→F only		F→P only		STEER	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-churn	RT	68.20	1023.2	**71.25**	259.9	71.06	230.2	71.06	230.2
		RT-J	72.26	1018.9	**74.59**	137.8	74.29	118.7	74.21	118.7
rel-amazon	user-churn	RT	64.14	1024.0	**65.41**	124.0	63.19	738.6	65.29	82.5
		RT-J	65.45	1021.3	65.30	105.5	63.50	742.2	**65.53**	68.1
rel-arxiv	paper-citation	RT	79.85	989.6	**80.30**	937.0	79.92	856.8	79.99	863.2
		RT-J	80.54	997.3	**80.78**	941.8	79.81	873.8	79.99	862.7
rel-avito	user-clicks	RT	51.92	613.6	57.31	545.9	61.56	497.0	**63.80**	487.5
		RT-J	**49.64**	612.2	42.50	549.6	43.51	476.0	40.52	463.6
rel-avito	user-visits	RT	55.98	706.7	58.95	418.4	56.08	676.9	**59.47**	374.8
		RT-J	63.77	731.7	64.44	210.3	64.05	703.5	**64.94**	150.7
rel-f1	driver-dnf	RT	78.55	1005.2	80.26	1002.3	80.30	1000.9	**81.40**	999.4
		RT-J	**78.50**	1001.5	77.37	1000.8	78.32	997.2	77.58	995.3
rel-f1	driver-top3	RT	89.26	1000.2	90.49	1000.2	89.72	1000.2	**90.70**	1000.2
		RT-J	89.23	996.6	**89.46**	996.5	89.45	996.6	89.15	996.8
rel-hm	user-churn	RT	63.80	1024.0	64.48	595.0	**65.75**	918.3	65.40	283.3
		RT-J	**63.63**	1020.3	61.50	543.0	61.61	910.4	61.49	197.6
rel-ratebeer	beer-rating-churn	RT	68.25	1024.0	67.19	1024.0	70.59	1022.7	**70.93**	1005.8
		RT-J	66.98	1011.7	71.05	1015.0	67.12	1009.5	**74.82**	976.7
rel-ratebeer	brewer-dormant	RT	**71.79**	1024.0	70.32	691.4	70.56	981.8	69.54	621.4
		RT-J	61.66	1007.8	62.08	695.7	**62.51**	965.5	61.83	631.8
rel-ratebeer	user-rating-churn	RT	73.38	1024.0	71.19	991.6	**74.31**	913.4	73.02	887.3
		RT-J	84.29	1013.1	**86.02**	991.6	85.82	913.5	85.72	898.2
rel-stack	user-badge	RT	**81.69**	355.6	78.94	256.8	79.14	303.2	75.89	134.2
		RT-J	89.53	357.8	89.65	256.3	**90.38**	311.8	87.01	125.0
rel-stack	user-engagement	RT	82.83	981.5	**85.70**	765.3	81.04	919.0	85.15	656.8
		RT-J	87.59	975.0	**88.41**	833.8	87.31	934.4	87.52	737.3
rel-trial	study-outcome	RT	50.05	1022.8	48.98	26.4	**53.53**	899.4	48.83	25.7
		RT-J	51.08	1011.0	52.53	25.8	48.05	756.6	**52.68**	25.2
Mean	(14 tasks)	RT	69.98	915.6	70.77	617.0	71.20	782.7	**71.46**	546.6
		RT-J	71.72	912.6	**71.83**	593.1	71.12	765.0	71.64	517.7
Δ	vs. baseline	RT	—	—	+0.79	-32.6%	+1.22	-14.5%	+1.48	-40.3%
		RT-J	—	—	+0.11	-35.0%	-0.60	-16.2%	-0.08	-43.3%

Table 8: Direction of traversal ablation: regression tasks (R^2 %, higher is better). Used in Figure 6. Ablation results use only the RT and RT-J models and the Claude LLM. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), P→F-only traversal, F→P-only traversal, and the STEER traversal filter that combines P→F and F→P. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 10 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		P→F only		F→P only		STEER	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-ltv	RT	**4.59**	957.2	4.46	954.1	4.26	203.5	4.26	203.5
		RT-J	59.31	951.5	58.89	945.4	60.00	108.4	**60.17**	108.4
rel-amazon	user-ltv	RT	15.20	1023.2	24.13	119.5	25.99	851.9	**26.26**	100.8
		RT-J	**61.52**	1021.3	53.86	105.5	53.66	602.3	57.19	86.9
rel-arxiv	author-publication	RT	22.77	740.5	**24.75**	491.4	22.64	680.6	24.75	491.4
		RT-J	18.46	728.8	30.50	453.3	21.46	684.0	**33.45**	474.4
rel-avito	ad-ctr	RT	4.19	1024.0	4.56	1024.0	**6.90**	1000.9	5.88	999.4
		RT-J	**2.44**	1020.8	-0.28	1020.2	2.09	995.2	-3.02	994.6
rel-f1	driver-position	RT	**45.31**	982.6	45.26	982.6	44.72	982.6	44.69	982.6
		RT-J	55.40	978.9	55.78	978.6	**55.84**	978.8	55.34	978.4
rel-hm	item-sales	RT	21.19	751.5	21.98	409.2	**23.74**	313.4	23.74	313.4
		RT-J	44.12	729.3	**44.49**	250.2	44.40	186.6	44.30	186.6
rel-ratebeer	user-rating-count	RT	12.91	1024.0	12.61	613.5	**27.15**	719.7	19.47	420.2
		RT-J	40.83	1013.0	38.75	613.1	**54.89**	706.6	45.06	272.6
rel-stack	post-votes	RT	28.75	1003.0	24.57	868.6	28.07	708.6	**29.18**	413.5
		RT-J	**86.08**	989.1	85.22	581.4	85.44	714.2	85.25	139.5
rel-trial	site-success	RT	14.78	1015.3	33.45	986.0	18.06	977.9	**34.83**	944.9
		RT-J	43.54	924.8	**47.24**	826.7	45.19	837.2	46.66	794.8
rel-trial	study-adverse	RT	**3.57**	1024.0	-0.82	21.9	3.44	879.0	-0.92	21.7
		RT-J	-7.37	1014.1	-11.90	22.2	**-6.84**	814.2	-12.19	22.0
Mean	(10 tasks)	RT	17.33	954.5	19.49	647.1	20.50	731.8	**21.21**	489.1
		RT-J	40.43	937.2	40.25	579.7	**41.61**	662.8	41.22	405.8
Δ	vs. baseline	RT	—	—	+2.17	-32.2%	+3.17	-23.3%	+3.89	-48.8%
		RT-J	—	—	-0.18	-38.1%	+1.18	-29.3%	+0.79	-56.7%

Table 9: Prompting strategy ablation: classification tasks (AUROC %, higher is better). Used in Table 1. Ablation results use only the RT and RT-J models. For each task and model, we report test accuracy with the baseline (random sampling). We also report simple prompting and STEER's smart prompting with four LLM variants: Claude and GPT with high and default reasoning effort. Within each simple/STEER pair, bold indicates the higher accuracy (ties favor STEER). The Mean block at the bottom of the table reports the per-model average across all 14 tasks. The Δ block reports the change relative to the baseline (point difference).

Database	Task	RFM	Base.	Claude High		Claude Default		GPT High		GPT Default	
				Simple	STEER	Simple	STEER	Simple	STEER	Simple	STEER
rel-amazon	item-churn	RT	68.20	**71.25**	71.06	**71.25**	71.07	71.06	**71.07**	71.06	**71.06**
		RT-J	72.26	**74.76**	74.21	**74.76**	74.41	74.34	**74.41**	**74.34**	74.21
rel-amazon	user-churn	RT	64.14	**65.41**	65.29	**65.41**	65.29	64.55	**64.55**	**65.41**	65.29
		RT-J	65.45	**65.54**	65.53	**65.54**	65.53	65.05	**65.05**	**65.54**	65.53
rel-arxiv	paper-citation	RT	79.85	79.85	**79.99**	**79.85**	79.83	74.21	**79.39**	74.21	**79.12**
		RT-J	80.54	**80.51**	79.99	**80.51**	80.27	75.01	**80.81**	75.01	**79.43**
rel-avito	user-clicks	RT	51.92	**64.79**	63.80	58.30	**64.60**	51.92	**63.02**	58.30	**58.36**
		RT-J	49.64	34.01	**40.52**	**41.76**	40.84	**45.69**	45.30	41.76	**44.42**
rel-avito	user-visits	RT	55.98	58.81	**59.47**	**58.97**	58.33	55.98	**59.57**	55.98	**57.14**
		RT-J	63.77	64.76	**64.94**	**64.78**	64.35	63.26	**65.07**	63.26	**65.53**
rel-f1	driver-dnf	RT	78.55	78.74	**81.40**	78.74	**80.64**	78.74	**80.62**	78.74	**79.96**
		RT-J	78.50	**79.19**	77.58	**79.19**	78.06	**79.19**	78.44	**79.19**	78.00
rel-f1	driver-top3	RT	89.26	89.36	**90.70**	89.68	**90.62**	**89.58**	89.35	89.48	**90.05**
		RT-J	89.23	**90.07**	89.15	**89.96**	89.18	**89.41**	89.18	**89.54**	89.06
rel-hm	user-churn	RT	63.80	64.48	**65.40**	64.48	**65.40**	65.26	**65.40**	64.48	**65.40**
		RT-J	63.63	61.30	**61.49**	61.30	**61.49**	60.96	**61.49**	61.30	**61.49**
rel-ratebeer	beer-rating-churn	RT	68.25	67.42	**70.93**	63.83	**68.36**	64.08	**70.25**	67.61	**69.98**
		RT-J	66.98	68.06	**74.82**	67.94	**71.77**	68.02	**75.75**	68.07	**74.15**
rel-ratebeer	brewer-dormant	RT	71.79	**70.07**	69.54	**70.86**	69.11	**70.01**	70.01	**70.07**	69.12
		RT-J	61.66	**62.52**	61.83	**62.49**	61.39	62.59	**63.96**	**62.52**	62.28
rel-ratebeer	user-rating-churn	RT	73.38	**73.32**	73.02	73.34	**73.82**	**81.51**	72.37	**81.51**	78.40
		RT-J	84.29	84.22	**85.72**	83.64	**84.56**	84.07	**86.10**	**84.07**	81.86
rel-stack	user-badge	RT	81.69	**81.49**	75.89	**81.69**	75.89	70.58	**78.77**	71.35	**77.60**
		RT-J	89.53	**89.78**	87.01	86.52	**87.01**	83.62	**87.50**	84.27	**87.82**
rel-stack	user-engagement	RT	82.83	82.83	**85.15**	82.57	**85.00**	**88.39**	85.80	83.22	**87.47**
		RT-J	87.59	87.03	**87.52**	**90.59**	90.29	**87.42**	86.81	**92.12**	90.44
rel-trial	study-outcome	RT	50.05	**49.15**	48.83	49.15	**50.24**	49.15	**49.99**	**49.15**	48.83
		RT-J	51.08	49.83	**52.68**	49.83	**52.08**	49.83	**53.26**	49.83	**52.68**
Mean	(14 tasks)	RT	69.98	71.21	**71.46**	70.58	**71.30**	69.64	**71.44**	70.04	**71.27**
		RT-J	71.72	70.83	**71.64**	71.34	**71.52**	70.60	**72.37**	70.77	**71.92**
Δ	vs. baseline	RT	—	+1.23	+1.48	+0.60	+1.32	-0.33	+1.46	+0.06	+1.29
		RT-J	—	-0.90	-0.08	-0.38	-0.21	-1.12	+0.64	-0.95	+0.20

Table 10: Prompting strategy ablation: regression tasks (R^2 %, higher is better). Used in Table 1. Ablation results use only the RT and RT-J models. For each task and model, we report test accuracy with the baseline (random sampling). We also report simple prompting and STEER's smart prompting with four LLM variants: Claude and GPT with high and default reasoning effort. Within each simple/STEER pair, bold indicates the higher accuracy (ties favor STEER). The Mean block at the bottom of the table reports the per-model average across all 10 tasks. The Δ block reports the change relative to the baseline (point difference).

Database	Task	RFM	Base.	Claude High		Claude Default		GPT High		GPT Default	
				Simple	STEER	Simple	STEER	Simple	STEER	Simple	STEER
rel-amazon	item-ltv	RT	4.59	4.00	**4.26**	4.00	**4.26**	4.00	**4.26**	**4.26**	4.22
		RT-J	59.31	59.05	**60.17**	59.05	**60.17**	59.05	**60.17**	**60.17**	59.84
rel-amazon	user-ltv	RT	15.20	24.13	**26.26**	24.13	**25.29**	**25.73**	25.29	**25.73**	23.90
		RT-J	61.52	54.60	**57.19**	54.60	**56.01**	54.41	**56.01**	54.41	**55.79**
rel-arxiv	author-publication	RT	22.77	22.99	**24.75**	22.80	**24.30**	23.14	**23.99**	22.57	**23.15**
		RT-J	18.46	14.56	**33.45**	14.57	**31.08**	14.64	**35.36**	14.76	**26.00**
rel-avito	ad-ctr	RT	4.19	**8.07**	5.88	**8.07**	3.82	4.48	**7.70**	4.48	**6.03**
		RT-J	2.44	**1.84**	-3.02	**1.84**	-3.40	**1.87**	-2.77	**1.87**	-1.92
rel-f1	driver-position	RT	45.31	**45.14**	44.69	45.38	**45.72**	45.33	**46.90**	45.18	**45.25**
		RT-J	55.40	**56.29**	55.34	**57.24**	55.67	**56.63**	55.55	**56.78**	55.35
rel-hm	item-sales	RT	21.19	21.98	**23.74**	21.98	**23.74**	**21.98**	21.62	**21.98**	21.22
		RT-J	44.12	**44.42**	44.30	**44.42**	44.30	**44.42**	42.81	**44.42**	43.49
rel-ratebeer	user-rating-count	RT	12.91	15.41	**19.47**	20.66	**26.18**	8.51	**18.21**	22.47	**26.10**
		RT-J	40.83	32.44	**45.06**	32.33	**44.84**	30.91	**44.66**	**46.40**	44.90
rel-stack	post-votes	RT	28.75	27.31	**29.18**	**27.31**	19.04	23.05	**28.22**	23.05	**27.74**
		RT-J	86.08	**85.68**	85.25	**85.68**	84.83	**84.59**	81.05	**84.59**	84.09
rel-trial	site-success	RT	14.78	1.81	**34.83**	1.84	**37.40**	2.20	**22.89**	-2.59	**22.49**
		RT-J	43.54	-0.69	**46.66**	-0.82	**44.84**	-1.10	**33.44**	-0.59	**33.30**
rel-trial	study-adverse	RT	3.57	**5.05**	-0.92	4.30	**5.37**	**5.05**	3.97	4.30	**5.86**
		RT-J	-7.37	**-6.97**	-12.19	**-6.75**	-7.99	**-6.97**	-7.18	**-6.75**	-8.07
Mean	(10 tasks)	RT	17.33	17.59	**21.21**	18.05	**21.51**	16.35	**20.31**	17.14	**20.60**
		RT-J	40.43	34.12	**41.22**	34.22	**41.03**	33.85	**39.91**	35.61	**39.28**
Δ	vs. baseline	RT	—	+0.26	+3.89	+0.72	+4.18	-0.98	+2.98	-0.18	+3.27
		RT-J	—	-6.31	+0.79	-6.22	+0.60	-6.59	-0.52	-4.83	-1.16

Table 11: Column filter ablation: classification tasks (AUROC %, higher is better). Used in Table 2. Ablation results use only the RT and RT-J models and the Claude LLM. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), STEER without column filter, and STEER with column filter. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 14 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		STEER		STEER + Col. Filter	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-churn	RT	68.20	1023.2	**71.06**	230.2	69.98	91.3
		RT-J	72.26	1018.9	**74.21**	118.7	72.67	50.8
rel-amazon	user-churn	RT	64.14	1024.0	**65.29**	82.5	63.75	35.5
		RT-J	65.45	1021.3	**65.53**	68.1	65.38	30.2
rel-arxiv	paper-citation	RT	79.85	989.6	**79.99**	863.2	78.37	862.0
		RT-J	**80.54**	997.3	79.99	862.7	79.34	862.8
rel-avito	user-clicks	RT	51.92	613.6	**63.80**	487.5	61.20	482.2
		RT-J	**49.64**	612.2	40.52	463.6	43.74	458.8
rel-avito	user-visits	RT	55.98	706.7	**59.47**	374.8	56.55	201.5
		RT-J	63.77	731.7	**64.94**	150.7	62.47	59.3
rel-f1	driver-dnf	RT	78.55	1005.2	**81.40**	999.4	81.27	999.3
		RT-J	78.50	1001.5	77.58	995.3	**80.32**	997.6
rel-f1	driver-top3	RT	89.26	1000.2	**90.70**	1000.2	90.52	1000.1
		RT-J	89.23	996.6	89.15	996.8	**89.75**	999.2
rel-hm	user-churn	RT	63.80	1024.0	65.40	283.3	**65.77**	62.4
		RT-J	**63.63**	1020.3	61.49	197.6	62.33	43.0
rel-ratebeer	beer-rating-churn	RT	68.25	1024.0	70.93	1005.8	**71.71**	1001.5
		RT-J	66.98	1011.7	74.82	976.7	**84.53**	972.7
rel-ratebeer	brewer-dormant	RT	**71.79**	1024.0	69.54	621.4	71.19	142.9
		RT-J	61.66	1007.8	**61.83**	631.8	60.30	78.9
rel-ratebeer	user-rating-churn	RT	73.38	1024.0	73.02	887.3	**79.56**	883.9
		RT-J	84.29	1013.1	85.72	898.2	**89.04**	902.3
rel-stack	user-badge	RT	**81.69**	355.6	75.89	134.2	72.24	126.3
		RT-J	**89.53**	357.8	87.01	125.0	80.93	110.5
rel-stack	user-engagement	RT	82.83	981.5	**85.15**	656.8	78.09	646.8
		RT-J	**87.59**	975.0	87.52	737.3	86.09	735.0
rel-trial	study-outcome	RT	50.05	1022.8	48.83	25.7	**51.33**	8.2
		RT-J	51.08	1011.0	52.68	25.2	**53.91**	7.7
Mean	(14 tasks)	RT	69.98	915.6	**71.46**	546.6	70.82	467.4
		RT-J	71.72	912.6	71.64	517.7	**72.20**	450.6
Δ	vs. baseline	RT	—	—	+1.48	-40.3%	+0.84	-48.9%
		RT-J	—	—	-0.08	-43.3%	+0.48	-50.6%

Table 12: Column filter ablation: regression tasks (R^2 %, higher is better). Used in Table 2. Ablation results use only the RT and RT-J models and the Claude LLM. For each task and model, we report test accuracy and number of tokens with the baseline (random sampling), STEER without column filter, and STEER with column filter. Within each row, bold indicates the highest test score and the highlighted cell indicates the lowest token count (ties are broken in favor of one non-baseline condition). The Mean block at the bottom of the table reports the per-model average across all 10 tasks. The Δ block reports the change relative to the baseline (point difference for test accuracy, and percentage reduction for tokens).

Database	Task	RFM	Baseline		STEER		STEER + Col. Filter	
			Test ↑	Tokens ↓	Test ↑	Tokens ↓	Test ↑	Tokens ↓
rel-amazon	item-ltv	RT	4.59	957.2	4.26	203.5	**33.25**	42.9
		RT-J	59.31	951.5	**60.17**	108.4	59.64	30.2
rel-amazon	user-ltv	RT	15.20	1023.2	26.26	100.8	**26.95**	50.1
		RT-J	**61.52**	1021.3	57.19	86.9	22.33	41.7
rel-arxiv	author-publication	RT	22.77	740.5	**24.75**	491.4	11.63	480.0
		RT-J	18.46	728.8	**33.45**	474.4	19.43	461.4
rel-avito	ad-ctr	RT	4.19	1024.0	5.88	999.4	**7.02**	999.3
		RT-J	**2.44**	1020.8	-3.02	994.6	-6.72	995.4
rel-f1	driver-position	RT	**45.31**	982.6	44.69	982.6	44.07	982.5
		RT-J	**55.40**	978.9	55.34	978.4	54.59	980.6
rel-hm	item-sales	RT	21.19	751.5	**23.74**	313.4	23.34	245.5
		RT-J	44.12	729.3	**44.30**	186.6	40.63	142.3
rel-ratebeer	user-rating-count	RT	12.91	1024.0	19.47	420.2	**50.34**	30.3
		RT-J	40.83	1013.0	45.06	272.6	**52.65**	22.3
rel-stack	post-votes	RT	28.75	1003.0	29.18	413.5	**29.80**	392.2
		RT-J	**86.08**	989.1	85.25	139.5	82.42	98.0
rel-trial	site-success	RT	14.78	1015.3	34.83	944.9	**42.74**	864.6
		RT-J	43.54	924.8	46.66	794.8	**63.62**	762.1
rel-trial	study-adverse	RT	**3.57**	1024.0	-0.92	21.7	-3.13	5.0
		RT-J	**-7.37**	1014.1	-12.19	22.0	-7.48	5.0
Mean	(10 tasks)	RT	17.33	954.5	21.21	489.1	**26.60**	409.2
		RT-J	40.43	937.2	**41.22**	405.8	38.11	353.9
Δ	vs. baseline	RT	—	—	+3.89	-48.8%	+9.28	-57.1%
		RT-J	—	—	+0.79	-56.7%	-2.32	-62.2%