WHY TABULAR FOUNDATION MODELS EXCEL ON RELATIONAL DATA THEY HAVE NEVER SEEN

Anonymous authors
Paper under double-blind review

ABSTRACT

Given a relational database (RDB) storing heterogeneous tabular information, how can we predict missing (or future) values in some target column of interest? As candidate targets are numerous in enterprise settings, it is preferable to avoid learning a new model from scratch each time there is a new prediction task. Frozen foundation models (FMs) based on in-context learning (ICL) and RDB-specific pre-training provide a viable solution, but the ideal design still remains largely unresolved. Perhaps unexpectedly, preliminary evidence has recently shown that existing *single-table FMs* (aka TFMs), when combined with suitable parameter-free relational encoders, can be quite competitive in practice despite *zero exposure to RDB pre-training or relational data of any kind*. In this work we first empirically confirm this phenomenon with a more diverse array of experiments. We then turn to our main contribution, namely, surfacing multiple complementary explanations for why TFMs behave so effectively. These explanations ultimately hinge on innate synergies whereby parameter-free encoders map relational effects to fixed-length embeddings that loosely mimic distributions for which TFMs were originally designed; in so doing, desirable single-table properties transfer to relational settings (e.g., convergent posterior predictive distributions). Collectively, our analysis forges new pathways for improving RDB FMs, particularly when we prioritize interpretability and inevitable integration within agentic systems.

1 INTRODUCTION

Relational databases (RDBs), housing collections of interrelated tables, are indispensable across enterprise applications and e-commerce platforms (Garcia-Molina et al., 2009). From a machine learning standpoint, the information stored within just a single RDB often contains countless possibilities for predictive modeling, covering targets such as customer retention, user churn, or click-through rates (Dave et al., 2014; Motl & Schulte, 2015; Ni et al., 2019; Zykov et al., 2022). Of course retraining a new model from scratch to address each new predictive task of interest may be prohibitively laborious or expensive, which motivates a new class of foundation models (FMs) specifically tailored for RDBs. Among others, these vary from LLM-based approaches that digest serialized RDB content (Wydmuch et al., 2024), to specialized Transformers pre-trained based on in-context learning (ICL) (Fey et al., 2025; Hudovernik et al., 2026; Ranjan et al., 2026b; Wang et al., 2026b), the common theme in each case being a *frozen* architecture capable of making predictions involving arbitrary unseen RDBs.

As a robust and scalable alternative, it has recently been demonstrated that *single-table FMs* (aka TFMs), when combined with suitable parameter-free subgraph encoders for capturing multi-table relational context, can perform well without requiring any RDB-specific pre-training (Grinsztajn et al., 2026; Hayler et al., 2026; Xu et al., 2026). After presenting background RDB context in Section 2, our work herein centers around confirming and explaining this unexpected performance via:

- **Comprehensive demonstration.** We provide thorough empirical comparison of wide-varying RDB foundation models across diverse benchmarking tasks. These results from Section 3 solidify the surprising effectiveness of TFMs repurposed for relational data.
- **Explanatory theory.** In Sections 4–7, we elucidate multiple complementary factors that help explain this phenomenon and broader merits of TFMs. Specifically, we show that TFM-based approaches: (i) naturally handle varying degrees of relational context, forming a continuum from single-tables to rich RDB schemas; (ii) transfer desirable asymptotic properties to RDB use cases, helping to compensate for mismatches between pre-training prior and practical deployment distributions; (iii) are not actually limited by expressivity concerns, even with parameter-free encoders and no prior RDB exposure, and (iv) integrate particularly well with agentic systems.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

2 BASICS OF RDB FOUNDATION MODELS

The goal of an RDB foundation model is to retain applicability across unseen RDBs, ideally with no retraining for each new predictive task. Beyond direct LLM-prompting approaches (Wydmuch et al., 2024), we now discuss recent pipelines designed to natively handle RDB/tabular data.

2.1 PREDICTIVE MODELING WITH RELATIONAL CONTEXT

Suppose we have training data $\{\boldsymbol{x}_i, y_i\}_{i=1}^n = \{\boldsymbol{X}, \boldsymbol{y}\}$ composed of instance feature rows $\boldsymbol{x}_i \in \mathcal{X}^d$ along with associated instance labels $y_i \in \mathcal{Y}$ for all i. The objective is to learn a parameterized model f_θ that achieves $f_\theta(\boldsymbol{x}_{\text{test}}) \approx y_{\text{test}}$ at any test point $\{\boldsymbol{x}_{\text{test}}, y_{\text{test}}\}$, where in practice y_{test} is unknown. RDB predictive modeling generalizes this canonical ML scenario by integrating an additional set of auxiliary data tables $\mathcal{T} = \{\boldsymbol{T}^k\}_{k=1}^K$, where $\boldsymbol{T}^k \in \mathfrak{T}^{n_k \times d_k}$ denotes the k-th table associated with a given entity type. Each table row corresponds with a single instance of that entity (e.g., a purchased item), while columns reflect heterogeneous instance attributes (e.g., product information). Note that w.l.o.g., we also assume $\boldsymbol{T}^1 = \boldsymbol{X}$. Lastly, an RDB also includes a set of relations $\mathcal{R}$ expressed as primary key (PK), foreign key (FK) pairs (Garcia-Molina et al., 2009); these may be absorbed into the constituent tables of $\mathcal{T}$ as additional columns when convenient. Both supervised methods and FMs make use of these auxiliary tables for making predictions.

2.2 COMPUTING RDB EMBEDDINGS

To make use of relational context for predictive tasks, the basic strategy is to first convert a given labeled RDB $\mathcal{D} := \{\mathcal{T}, \mathcal{R}, \boldsymbol{y}\}$ to a heterogeneous graph $\mathcal{G}$. This typically involves treating each table $\boldsymbol{T}^k$ as a node type and each row i within a given $\boldsymbol{T}^k$ as a node with features $\boldsymbol{t}_i^k$ (Cvitkovic, 2020; Dwivedi et al., 2025; Fey et al., 2023; Zhang et al., 2023a;b). Directed edges can then be formed using each PK-FK pair within $\mathcal{R}$. We next sample H-hop subgraphs or ego-networks $\mathcal{G}_H(\boldsymbol{x}_i)$ from $\mathcal{G}$ that are centered at each target row $\boldsymbol{x}_i$ within $\boldsymbol{X}$. Importantly, sampling for *temporal* RDBs should exclude nodes with time-stamps later than $\boldsymbol{x}_i$. These subgraphs form the input to an RDB encoder g_ϕ, instantiated as some form of graph neural network or Transformer architecture with (optional) parameters ϕ. This facilitates computation of *fixed-length* embeddings $\boldsymbol{z}_i = g_\phi[\mathcal{G}_H(\boldsymbol{x}_i)] \in \mathbb{R}^{d_z}$.

2.3 COMBINING WITH RDB PREDICTION HEADS

Schema-agnostic models. Provided the encoder g_ϕ is designed to digest a broad spectrum of input RDB schema using a shared representational form, then it is possible to train a single composite model $q_\theta\big(y \mid \boldsymbol{z}_{\text{test}}\big) \equiv q_\theta\big(y \mid g_\phi\big[\mathcal{G}_H(\boldsymbol{x}_{\text{test}})\big]\big)$ to assign high probability to labels $y = y_{\text{test}}$ (or $y \approx y_{\text{test}}$ for regression tasks) across a range of real-world RDBs (Ranjan et al., 2026a; Wang et al., 2025; Wu et al., 2025). Here $\boldsymbol{x}_{\text{test}}$ can be viewed as an additional unlabeled row(s) appended to $\boldsymbol{X}$. Such models can in principle be applied to new unseen RDBs without further training.

Pre-trained ICL-based models. A limitation of the schema-agnostic models from above is that the effective context available for making predictions is merely the subgraph $\mathcal{G}_H(\boldsymbol{x}_{\text{test}})$, which in isolation may not contain sufficient information pertaining to broader labeling patterns across a vast RDB. In-context learning (ICL) provides a natural workaround by granting the prediction head diverse samples of multiple analogous subgraphs extracted from each new RDB of interest. Specifically, ICL-based architectures derive predictions from the revised distribution

$$q_\theta\big(y \mid \boldsymbol{z}_{\text{test}}, \mathcal{S}_n\big), \text{ with } \mathcal{S}_n = \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \ \boldsymbol{z} = g_\phi[\mathcal{G}_H(\boldsymbol{x})],$$

where $\boldsymbol{z}_{\text{test}}$ and each $\boldsymbol{z}_i$ are obtained by passing unlabeled and labeled rows of $\boldsymbol{X}$ through the shared encoder g_ϕ. It is also possible to apply g_ϕ to revised tabular embeddings $\boldsymbol{U}^k \leftarrow \boldsymbol{T}^k$ rather than original tables (Hudovernik et al., 2026). The prediction head q_θ is implemented via various Transformer architectures that grant flexibility over n and d_z. Finally, model pre-training is conducted using large volumes of synthetically-generated data, possibly combined with real-world RDBs (Chen et al., 2026; Fey et al., 2025; Hudovernik et al., 2026; Kothapalli et al., 2026; Ranjan et al., 2026b).

Retrofitting TFMs with a relational encoder. Given the growing range of single-table FMs (Hollmann et al., 2025; Jingang et al., 2025; Qu et al., 2026; Zhang et al., 2025b;a), it is natural to consider combining with a *parameter-free* encoder g_{pf}, for flattening RDB neighborhoods into fixed-length tabular features, such that the entire pipeline requires no RDB-specific training/pre-training (Grinsztajn et al., 2026; Hayler et al., 2026; Xu et al., 2026). Herein we aim to confirm and explain the effectiveness of these simple approaches relative to more complex alternatives.

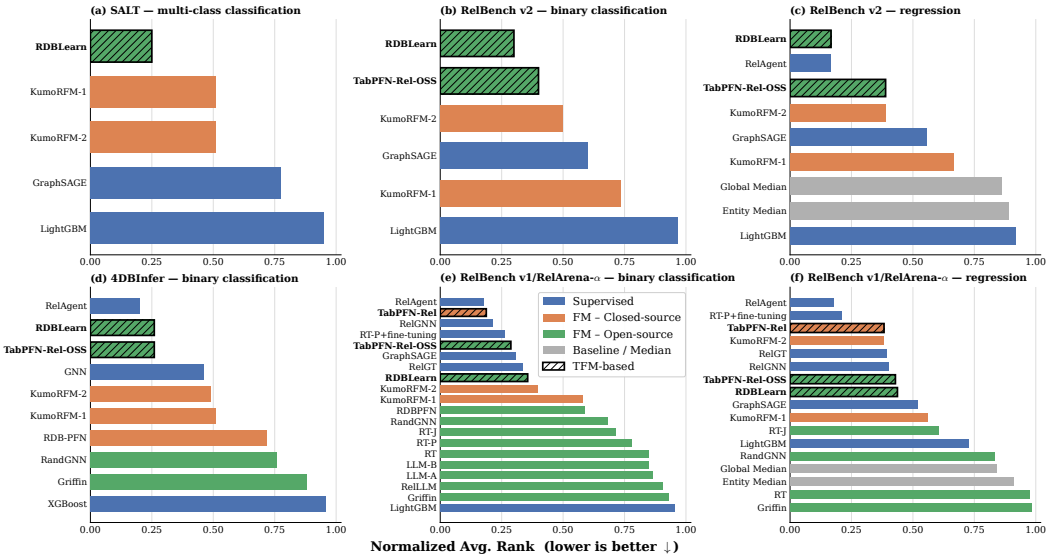


Figure 1: TFM-based approaches outperform all other RDB FMs on all benchmarks.

3 More Empirical Evidence that TFMs Excel on RDBs

Models. For TFM-based approaches, we adopt *RDBLearn* (Xu et al., 2026) and *TabPFN-Rel* (Hayler et al., 2026) based on analogous principles; TabPFN-Rel has both open- and closed-source versions (the latter including proprietary text handling and unknown RDB-related modifications). For RDB-native FMs, we include a variety of both open-source and closed-source options with published results based on author-specific tuning. Note that most models only provide results on a subset of available benchmarks. Lastly, we cover a handful of supervised baselines that represent either SOTA performance or classical methodology as a reference point. Appendix E includes all baseline model descriptions as well as further details of how results were obtained.

Results. We compare across a suite of six diverse benchmarks from Appendix E, each composed of multiple RDB predictive tasks. We remark that it is routine to only report results on RelBench-v1 datasets/tasks (Jäger et al., 2026; Hayler et al., 2026; Kothapalli et al., 2026; Ranjan et al., 2026a;b), which are also equivalent to recent RelArena-α task designations at the time of this writing (Hayler et al., 2026). But this poses greater risk of overfitting to limited benchmarking data, e.g., settings/architectures suitable for RelBench-v1 need not always translate to stronger performance elsewhere. For the expanded benchmark set considered here, results are shown in Figure 1 and Appendix E.4, where TFM-based approaches achieve the best performance among all FMs, including even several variants that may involve pre-training with RDBs structured similarly to the benchmarks. We stress that RDBLearn and TabPFN-Rel-OSS were run directly from open-source repositories for all benchmarks *with no task-specific adjustments and zero exposure to RDB data of any kind at any stage*, i.e., no chance for any implicit test-set alignment on public benchmarks.

Why are TFMs so competitive on RDBs? It is initially reasonable to assume that pre-training RDB-specific FMs using RDB data is essential for achieving strong inference-time performance on unseen relational data. So how can TFMs paired with some g_{pf} possibly perform well, as we have demonstrated in Figure 1. The high-level argument from Xu et al. (2026) is that prior to seeing ICL samples, a data-agnostic encoder is incapable of adjudicating column roles, which may vary from dataset to dataset (or even task to task within a given RDB). This phenomenon advantages RDB encoders that only compress vertically *within* high-dimensional columns (where shared units facilitate interpretable aggregation), not horizontally *across* columns, where relevance is largely indeterminate without adequate label information. Conditioned on this restriction to vertical (intra-column) compression, Xu et al. (2026) detail how g_{pf} can result in minimal loss of expressiveness, while naturally pairing with existing TFMs to form a complete RDB pipeline. While promising as an analytical starting point, there remains considerable ambiguity regarding the rationale and surprising empirical efficacy of this simple recipe. In the remainder of this work we take important disambiguating steps, beginning with a discussion of the blurry line between single tables and RDBs, where TFMs already work well on the former.

4 SINGLE TABLES AND RDBS ON A CONTINUUM

Inducing RDBs from single tables. Conventionally, TFMs and benchmarking (i.e., single-table regimes) are treated separately from their RDB counterparts. But as we now examine, infinitesimal adjustments to an initial single-table setup facilitate gradual transitioning to multi-table relational scenarios, making any strict partitioning somewhat arbitrary. Concretely, beginning with a single-table scenario, meaning $K = 1$ and features $\boldsymbol{T}^1 = \boldsymbol{X}$, we can trivially transform these data into two or more tables with arbitrarily-small, labeled cross-row relational dependencies as follows:

1. Select any categorical column of $\boldsymbol{T}^1$; w.l.o.g. assume B categories. If no categorical columns exist, then we simply chose a numerical column and quantize into B bins.
2. Form a new auxiliary table $\boldsymbol{T}^2$ with a single length-B column, where elements are the unique categories obtained from step 1 above. Treat the lone column within $\boldsymbol{T}^2$ as a PK, and the original categorical column within $\boldsymbol{T}^1$ as an FK.
3. To induce more complex schema, we may optionally repeat steps 1–2 using other $\boldsymbol{T}^1$ columns.
4. Form subgraphs $\mathcal{G}_H(\boldsymbol{x})$ seeded at any row $\boldsymbol{x}$; provided $H \geq 2$, these neighborhoods will include multiple nodes associated with different rows of $\boldsymbol{X}$. This is because the first hop reaches an induced $\boldsymbol{T}^k$ ($k > 1$), while any second hop returns to $\boldsymbol{T}^1$. In doing so, cross-row relational dependencies are introduced via the edges within $\mathcal{G}_H(\boldsymbol{x})$.
5. Define the label generation function
$$y = \psi\big[(1-\tau) f_1(\boldsymbol{x}, \varepsilon_1) + \tau f_2\big(\mathcal{G}_H(\boldsymbol{x}) \backslash \boldsymbol{x}, \varepsilon_2\big)\big], \quad (1)$$
where $\tau \in [0, 1]$ represents a continuation parameter and $\{\varepsilon_1, \varepsilon_2\} \in \mathbb{R}^2$ are exogenous noise factors. Meanwhile, f_1 and f_2 are functions mapping their respective domains to a shared r-dimensional embedding space $\mathbb{R}^r$, after which $\psi : \mathbb{R}^r \to \mathcal{Y}$ converts to the final label space.

Continuum perspective. Any possible conditional distribution $p\left(y|\mathcal{G}_H(\boldsymbol{x})\right)$ can in principle be obtained by generating data via step 5, subject to mild regularity conditions (Peters et al., 2017)[Proposition 4.1]; likewise for special cases $p(y|\boldsymbol{x})$ (when $\tau = 0$) and $p\left(y|\mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}\right)$ (when $\tau = 1$). Hence given any set of tabular features, we can always generate a continuum of rich label distributions with varying degree of relational dependency, ranging from the canonical single-table regime ($\tau = 0$) to an interconnected multi-table structure ($\tau \neq 0$). We emphasize that our point here is not to replicate the most likely real-world processes by which RDBs are generated, nor discount the countless other ways single tables can be normalized into RDB-like schema. Instead, we are merely demonstrating that there is *no unambiguous line between single tables and multi-table relational data*, as relational effects can diversely form a continuum from non-existence to prominence.

Plausibility of TFMs on RDBs. TFMs by original design excel handling $\tau = 0$ cases, and some real-world RDBs likely approximate this regime (Akhter et al., 2026). Furthermore, if they are equipped with a suitably powerful encoder g_{pf}, such as the encoding strategy endorsed by Xu et al. (2026), then larger τ cases are also within reach as evidenced by Figure 1. Hence, TFMs are reasonable candidates for navigating relational continuums, providing a strong initial plausibility argument for their usage on RDBs. Conversely, if an alternative RDB FM does not work well on single-tables, what about performance on RDBs that are minimally different from single tables, and so on? Again, these considerations motivate RDB FMs that explicitly handle a broad relational spectrum as TFMs already do. In Appendix A.1 we closely examine limitations of adopting interpolation/ensembling between single-table FMs and RDB FMs as a simplistic workaround, proving a pessimistic bound (Proposition A.3) that occurs even with very minor modeling constraints. See also Appendix A.2 for broader benchmarking implications of the continuum perspective.

Empirical implications. We empirically probe implications of the continuum perspective, particularly the behavior of RDB FMs operating near single-table regimes. For this purpose, we first collect a set of tabular datasets (originally single table), whereby train/test splits are temporal as in typical RDB scenarios; see Appendix F for all setup details. We then induce RDBs using steps 1–4 from above; however, we retain the original dataset labels to maintain real-world authenticity. Note that these datasets, designed to test non-IID temporal shifts, may already exhibit small degrees of cross-row dependency even without artificially inducing such dependencies via (1). We compare TFM-based RDBLearn against the top RDB FMs (non-TFM-based) KumoRFM-2 (Hudovernik et al., 2026) and RT-J (Ranjan et al., 2026b) in Figure 2 and Appendix F. Overall, RDBLearn shows the best performance among RDB FMs. Meanwhile, although RT-J struggles with these particular tasks,

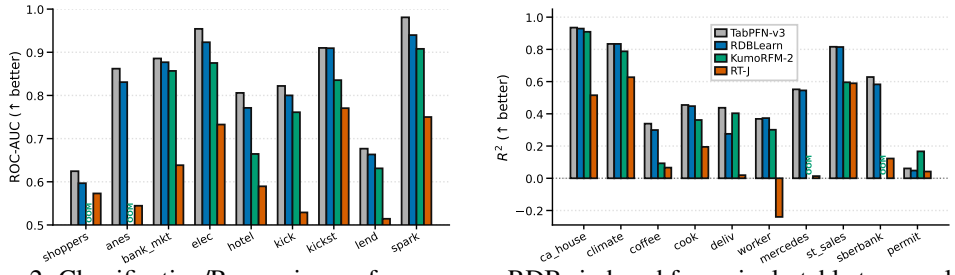


Figure 2: Classification/Regression performance on RDBs induced from single-table temporal data.

KumoRFM-2 performs significantly better (excluding 4 OOM cases). This is perhaps not surprising given that KumoRFM-2 pre-training explicitly transitions from single tables to relational data over an undisclosed multi-stage process (Hudovernik et al., 2026). Moreover, inspection of the KumoRFM-2 implementation reveals heavy reliance on the TabICL-v2 TFM architecture (Qu et al., 2026). Pure TFM-based methods like RDBLearn and TabPFN-Rel simply take these observations one step further by dispensing with potentially-complex RDB-specific pre-training altogether. Either way, it is our view that Figure 2-style benchmarking represents an informative hurdle for RDB FMs to pass.

5 Locality Considerations Favor TFM Usage

We now turn to an RDB generative process that facilitates the study of locality-based necessary conditions for removing bias. We then find that TFMs, with uniformly drawn ICL samples produced by a suitable g_{pf}, allow us to approximately satisfy these conditions with RDB data.

5.1 Necessary Conditions for Unbiased RDB FMs

Definition 5.1. (*continual RDB generation*) Let $\mathcal{P}$ denote a set of RDB distribution sequences $\rho := \{p_n\}_{n=m}^{\infty}$, while $\mathcal{D}_n \sim p_n$ is an RDB associated with ρ for any $n > m \geq 0$. We stipulate that all such $\mathcal{D}_n$ include a table $\boldsymbol{T}^1$ composed of labeled rows $\{\boldsymbol{x}_i, y_i\}_{i=0}^{n}$. These labels are assumed to follow the conditional $p\big(y_i|\mathcal{G}_H(\boldsymbol{x}_i)\big)$ that does not vary across n, but may vary with different $\rho \in \mathcal{P}$ (in contrast, we do not prevent the marginal $p\big(\mathcal{G}_H(\boldsymbol{x}_i)\big)$ from varying across n and ρ). We treat $\{\mathcal{G}_H(\boldsymbol{x}_0), y_0\} \equiv \{\mathcal{G}_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}\}$ as a test sample that remains fixed across n.

By design, Definition 5.1 allows us to consider distributions of RDBs whereby the overall size and number of ICL samples is growing indefinitely (these may extend into the past if timestamps are involved), but with a fixed evaluation point $\{\mathcal{G}_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}\}$ (akin to typical IID tabular cases). And even if real-world RDBs are not generated in exactly this way, with modest assumptions, the same RDB distribution can often be defined/generated in either a forward or backward fashion. Hence we preserve flexibility with this construction while facilitating asymptotic analyses as follows.

Consider a scenario where, as n becomes large, a given RDB FM achieves two desirable conditions: (i) the variance of the predicted label vanishes, and (ii) the bias relative to the true label distribution, conditioned on an encoder output, tends to zero. Mathematically, these conditions manifest as

$$\begin{aligned} q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big) - \mathbb{E}_{\mathcal{D}_n \sim p_n}\left[q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right] &\to 0 \quad \text{(variance)} \\ p\big(y_{\text{test}} \mid \boldsymbol{z}_{\text{test}}\big) - \mathbb{E}_{\mathcal{D}_n \sim p_n}\left[q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right] &\to 0 \quad \text{(bias)} \end{aligned} \tag{2}$$

almost surely as as $n \to \infty$. We are assuming here that $\boldsymbol{z} = g_\phi\big(\mathcal{G}_H(\boldsymbol{x})\big)$ and that $\{p_n\}_{n=m}^{\infty}$ represents a sequence of RDB distributions adhering to the generative process from Definition 5.1; likewise for $p\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big)$ which does not vary with n. Building on the rich conceptual framework from Nagler (2023), originally designed to analyze IID scenarios over single tables, we now examine *necessary* conditions on a given $\{q_\theta, g_\phi\}$ for achieving (2), noting that g_{pf} is a special case of g_ϕ.

Proposition 5.2. *Given an RDB FM defined by some $\{q_\theta, g_\phi\}$, and every $\rho \in \mathcal{P}$ adhering to Definition 5.1, assume (2) holds and that $p\big(y_i|\mathcal{G}_H(\boldsymbol{x}_i)\big) = p\big(y_i|\boldsymbol{z}_i\big)\ \forall i = 0, \ldots, n$. Then for any $\check{\rho} \in \mathcal{P}$ (and associated label generator $\check{p}$) there exists a sequence $\varepsilon_n \to 0$ such that almost surely*

$$\left|q_\theta\big(y_{test} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{test}\big) - q_\theta\big(y_{test} \mid \{\check{\boldsymbol{z}}_i, \check{y}_i\}_{i=1}^n, \boldsymbol{z}_{test}\big)\right| \xrightarrow{n\to\infty} 0, \tag{3}$$

where samples $\{\boldsymbol{z}_i, y_i\}_{i=1}^n$ are extracted from $\mathcal{D}_n \sim p_n$, samples $\{\check{\boldsymbol{z}}_i, \check{y}_i\}_{i=1}^n$ from $\check{\mathcal{D}}_n \sim \check{p}_n$, and $p\big(y \mid \boldsymbol{z}\big) = \check{p}\big(y \mid \boldsymbol{z}\big)$ if $\|\boldsymbol{z} - \boldsymbol{z}_{test}\| \leq \varepsilon_n$.

Intuitively, this result merely indicates that if we wish to obtain unbiased estimates of the label across broad distributional settings within $\mathcal{P}$ (as required by RDB foundation models in particular), then the output of our decoder function q_θ should depend only on ICL samples with encoder representations $\boldsymbol{z}_i$ that are *close* to the test representation $\boldsymbol{z}_{\text{test}}$. Note that

$$\{\boldsymbol{z}_i, y_i\}_{i=1}^n = \{\boldsymbol{z}_i, y_i\}_{\boldsymbol{z}_i \in \Omega_n} \cup \{\boldsymbol{z}_i, y_i\}_{\boldsymbol{z}_i \notin \Omega_n} \text{ with } \Omega_n = \{\boldsymbol{z} : \|\boldsymbol{z} - \boldsymbol{z}_{\text{test}}\| \leq \varepsilon_n\}, \tag{4}$$

while $\{\check{\boldsymbol{z}}_i, \check{y}_i\}_{i=1}^n$ decomposes analogously. When n grows large, the sets $\{\boldsymbol{z}_i, y_i\}_{\boldsymbol{z}_i \in \Omega_n}$ and $\{\check{\boldsymbol{z}}_i, \check{y}_i\}_{\check{\boldsymbol{z}}_i \in \Omega_n}$ become indistinguishable as Ω_n collapses around $\boldsymbol{z}_{\text{test}}$, hence they may influence the output of q_θ without conflicting with (3). In contrast, remaining samples not in Ω_n must have negligible influence for (3) to hold. This dependency on *locality* allows the FM to differentiate RDBs that have similar behavior around any given test point, but increasing separation elsewhere. We next discuss the implications for RDB FM design types, particularly those based on TFMs.

5.2 Natural Alignment with Single-Table FMs.

Uniform ICL sampling. When choosing ICL samples, we would like at least some associated with encoder representations $\{\boldsymbol{z}_i\}_{i=1}^n$ that are reasonably *close* to the test point representation $\boldsymbol{z}_{\text{test}}$, as required for unbiased estimation per Proposition 5.2. In the absence of domain-specific knowledge or strong a priori assumptions such as homophily,[1] we may not know in advance where these close samples will be located in a given RDB. And yet FM architectures should ideally address wide-ranging conditions with a shared data-agnostic strategy. Along these lines, uniform sampling (from rows of $\boldsymbol{X}$ with historical timestamps) is reasonable for approximating the full RDB context across generic new tasks, leading to at least some coverage of any new $\boldsymbol{z}_{\text{test}}$ with select elements of $\{\boldsymbol{z}_i\}_{i=1}^n$.[2] Moreover, in a sufficiently large RDB these uniform samples may even behave roughly as though they were IID, mirroring the single-table distributions for which TFMs were originally pre-trained. This alignment can play a critical role in achieving more robust performance as the number of ICL samples grows, which we will investigate next through the lens of *posterior predictive distributions* (PPDs) (Müller et al., 2021). The PPD serves a critical role, representing the target distribution that ICL-based TFMs attempt to approximate during pre-training (Hollmann et al., 2022).

Leveraging Single-Table PPDs. Given some RDB $\mathcal{D}$ drawn from ground-truth distribution p_{gt}, suppose we extract samples $\{\mathcal{G}_H(\boldsymbol{x}_i), y_i\}_{i=1}^n$ and test point $\{\mathcal{G}_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}\} = \{\mathcal{G}_H(\boldsymbol{x}_0), y_0\}$ satisfying

$$p_{\text{gt}}\big(\{\mathcal{G}_H(\boldsymbol{x}_i), y_i\}_{i=0}^n\big) \approx \prod_{i=0}^n p_{\text{gt}}\big(\mathcal{G}_H(\boldsymbol{x}_i), y_i\big) \text{ and } p_{\text{gt}}\big(y_i|\mathcal{G}_H(\boldsymbol{x}_i)\big) \approx p_{\text{gt}}\big(y_i|\boldsymbol{z}_i\big) \ \forall i. \tag{5}$$

These may be reasonable approximations if (i) uniform sampling is conducted over a sufficiently large RDB, and (ii) the embeddings $\boldsymbol{z}_i$ adequately reflect the dominant sources of discriminative information within H-hop neighborhoods. Granting the assumptions from (5) hold as n grows large, we may directly borrow an existing prior μ over *single-table* distributions $q \in \mathcal{Q}$ satisfying mild regularity conditions (Nagler, 2023). We then form a PPD that converges towards the ground-truth as

$$p_{\text{ppd}}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}, \mathcal{S}_n\big) \;=\; \int q\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big) d\mu\big(q|\mathcal{S}_n\big) \;\xrightarrow{n\to\infty}\; \hat{p}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big) \;\approx\; p_{\text{gt}}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big), \tag{6}$$

where $\hat{p}$ is the KL-optimal approximation of p_{gt} in $\mathcal{Q}$. This drift towards the ground-truth p_{gt} is essential for the following reason. While the distribution of pre-training data for TFMs may loosely resemble trends in real-world tables, there is certainly no exact correspondence. However, a perfect match with the prior is *not* required to achieve strong performance, provided that: (i) the prior assigns at least some nonzero probability mass to a $\hat{p}$ reasonably near p_{gt}, and (ii) we have collected enough uniformly-distributed ICL samples such that the true PPD narrows as $p_{\text{ppd}}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}, \mathcal{S}_n\big) \approx \hat{p}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big)$. This phenomenon (originally motivated from single-tables) may at least partially explain the observed empirical success generalizing to RDBs, whereby conditions mirroring the single-table case have been explicitly engineered through uniform sampling and encoders g_{pf} that do not mix across columns, preserving column identity and roles akin to typical tabular regimes (Xu et al., 2026). The latter plausibly increases the chances that condition (i) from above is roughly satisfied in regions around $\boldsymbol{z}_{\text{test}}$. In so doing we effectively align RDB neighborhoods with attributes of input data previously applied to successful TFMs, transferring similar efficacy.

[1] In the context of heterogeneous graphs or subgraphs defining RDB relations, *homophily* refers to the tendency of neighboring nodes of the same node type (not necessarily 1-hop neighbors) to share the same label, while *heterophily* represents the converse.

[2] By uniform sampling, we do not exclude temporal bias. For example, if we wish to prioritize recency, we can still uniformly sample across all RDB neighborhoods satisfying a time difference constraint.

5.3 Limitations of Alternative Possibilities

Bias from non-uniform sampling. The distance between local RDB neighborhoods, as arbitrated by encoder representations $\{\boldsymbol{z}_i\}_{i=1}^n$, determines the notion of locality upon which Proposition 5.2 is based. But it is nonetheless useful to contrast this notion with distinct options adopted by other types of RDB FMs. For example, the latest SOTA relational Transformer instantiation RT-J (Ranjan et al., 2026b), based on RDB-specific pre-training with real-world RDB data, involves forming ICL samples as the set of target rows in $\boldsymbol{T}^1$ that are most frequently visited over the course of a large number of fixed-length *random walks* seeded from $\boldsymbol{x}_{\text{test}}$. In this way we collect samples that are close to $\boldsymbol{x}_{\text{test}}$ in the sense of graph distance, *but not necessarily in terms of encoder representation*. For networks dominated by homophily effects, meaning nearby nodes in the graph have similar labels, this likely works well. But more broadly, within an RDB influenced by heterophily relationships, disconnected regions, or other spurious network effects, similar samples/nodes drawn from random walks may *not* be similar in terms of encoder representations, in which case the necessary conditions of Proposition 5.2 would not hold. This may also partially explain the performance of RT-J in Figure 2, where the infused relational structure is minimally related to the labels by experimental design.

Implications for the PPD. The distinction between uniform and non-uniform sampling also influences the PPD upon which ICL-based FMs depend. For example, regardless of whether the measure μ reflects a single-table or a full RDB pre-training prior, if random-walk sampling is adopted to obtain $\mathcal{S}_n$, then the core IID assumption from (5) used to establish (6) will no longer generally hold. Consequently, increasing the number of ICL samples need not approximate a desirable asymptotic regime, and the resulting PPD may bear little resemblance to the target $p_{\text{gt}}\big(y_{\text{test}}|\boldsymbol{z}_{\text{test}}\big)$. See Appendix C for a simple example whereby random-walks stall around the test node, failing to capture samples that lie elsewhere in the RDB as required to produce a meaningful PPD (the core issue is that $\mathcal{G}$ is disconnected). Hence instead of having n approximately IID samples from some $p_{\text{gt}}(\mathcal{G}_H(\boldsymbol{x}), y)$ (as with uniform sampling matched with a single-table prior μ), with random walk sampling we instead have just *one sample* from a much more complex, entangled joint distribution $p_{\text{gt}}\big(\{\mathcal{G}_H(\boldsymbol{x}_0), y_0, \}, \ldots, \{\mathcal{G}_H(\boldsymbol{x}_n), y_n\}\big)$. This dramatically increases the necessity of an accurate *RDB-specific* pre-training data prior to resolve ambiguity in the resulting PPD, as (6) need not hold any longer, even approximately. But such priors raise a new issue.

Dilemma of RDB-specific pre-training priors. Now suppose that μ is no longer simply a prior over distributions of single-tables with IID rows (as assumed by TFMs), but instead over distributions of full RDBs as adopted by multiple RDB FMs (Chen et al., 2026; Fey et al., 2025; Hudovernik et al., 2026; Ranjan et al., 2026b). But designing an accurate prior of this type is more challenging than in the single-table case, simply because of all the new relational dependencies introduced over multiple tables. This implies that misspecifications may actually be *greater* relative to real-world RDBs, and yet we cannot necessarily compensate with a larger n (as in TFM cases) since this need no longer induce an IID regime or the equivalent of (6). These observations suggest that, outside of specific domains where developing targeted pre-training priors may be more feasible, generic TFM-based pipelines still hold some sway. Plausibly, loose IID approximations with even modest compensatory effects as n grows large might often generalize better in new domains than highly-complex, entangled priors that do not factorize over ICL sample sets $\big(\{\mathcal{G}_H(\boldsymbol{x}_0), y_0, \}, \ldots, \{\mathcal{G}_H(\boldsymbol{x}_n), y_n\}\big)$.

6 Adequate Encoder Expressivity even with No parameters

We next address possible expressivity limitations associated with g_{pf} usage, the key ingredient which allows TFMs to avoid any RDB-specific training or pre-training. A natural presumption here is that parameter-free encoders place a potentially-burdensome restriction on the types of RDB neighborhood representations available for constructing ICL samples. We analyze this concern through a case study involving interaction terms, where we find that the added expressivity of parameterized encoders g_ϕ need not provide an unequivocal advantage over the much simpler g_{pf} adopted by TFMs.

Interaction terms as a case study. Prior work has noted that parameter-free encoders that process RDB columns individually (as mentioned in Section 3 and adopted by RDBLearn/TabPFN-Rel) are incapable of handling interaction terms between columns in auxiliary tables. Along these lines, consider the following representative example of this ostensible limitation discussed by Hudovernik et al. (2026); Huang et al. (2026). Suppose we are given and RDB with only two tables: A labeled target table $\boldsymbol{T}^1 = \boldsymbol{y}$ and a child table $\boldsymbol{T}^2 = [\boldsymbol{a}, \boldsymbol{b}]$ composed of two binary feature columns $\boldsymbol{a}$ and $\boldsymbol{b}$. A PK associated with each row of $\boldsymbol{T}^1$ is linked by FKs to one or more rows of $\boldsymbol{T}^2$, where $\mathcal{R}_i$ denotes

the set of rows within $\boldsymbol{T}^2$ that have FKs pointing to the PK associated with row i in $\boldsymbol{T}^1$. Ground-truth labels are then constructed to satisfy

$$y_i = \max_{i \in \mathcal{R}_i} [a_i \wedge b_i] \quad \forall i. \tag{7}$$

Clearly the element-wise interactions $a_i \wedge b_i$ determine the label, and so the argument goes, excluding such interactions is disproportionately problematic within any g_{pf} that encodes columns separately as advocated by Xu et al. (2026). We provide three counter-points to this concern:

***All* lossy encoding involves negative counter-examples.** The labels computed via (7) involve a binary function over $2|\mathcal{R}_i|$ possible input configurations. *But this is just one possible function sharing this input range*. Assuming no permutation invariance across rows, there are $4|\mathcal{R}_i|$ distinct binary functions expressible over the same input space (for each unique input, we can assign either $y_i = 0$ or $y_i = 1$, doubling $2|\mathcal{R}_i|$ to achieve the stated total). Without external knowledge mandating which function is correct, an encoder that fails to losslessly preserve the entire input range will make it impossible to guarantee that a subsequent decoder can determine the correct labeling rule; a single missing bit can flip any labeling rule leading to errors within an arbitrary downstream classifier.

More broadly, any possible lossy encoding, *whether parameterized or not*, opens up a set of inevitable counter-examples, whereby learning the correct function is infeasible; see Xu et al. (2026)[Proposition 4.2] as one representative scenario. This remains true even if task-specific labels are included within $\mathcal{G}_H(\boldsymbol{x})$; more on this in Appendix B. Consequently, the primary differentiating consideration is not whether counter-examples exist; rather, what matters is the relative frequency with which encoder-specific counter-examples appear in real-world RDBs.[3] In general, functional expressivity is likely more critical in *supervised* scenarios, where trainable degrees-of-freedom can be mapped to task-specific encoders that mitigate task-specific counter-examples (e.g., a particular set of interaction terms). But for ICL-based FMs restricted to a fixed data-agnostic form, other factors such as interpretability or sample complexity may matter more than parameterized expressivity per se.

Arbitrary interactions challenge *any* fixed encoder. Even if we restrict label generation to functions of interaction terms similar to (7), parameterized encoders still do not have a clear-cut advantage in real-world settings. This is because real-world auxiliary tables often have many columns. Excluding all higher-order interactions, there will still exist a quadratically expanding number of candidate interaction terms as the number of columns grow. Without knowing in advance which are important, lossy encoder compression, *parameterized or not*, will necessarily lead to prediction errors (i.e., there will always be cases where the interaction signal depends on missing information). Appendix B analyzes how historical labels within $\mathcal{G}_H(\boldsymbol{x})$ cannot generally resolve this ambiguity.

Parameter-free encoders can (sometimes) handle interactions. For arguments sake, the presumption thus far in this section has been that a column-wise parameter-free g_{pf} cannot represent interaction terms. But actually, this is not strictly true for the following two reasons. First, models like RDBLearn or TabPFN-Rel can directly handle interactions between compressed representations of individual columns; these representations form additional columns within $\boldsymbol{Z} \in \mathbb{R}^{n \times d_z}$, and TFMs can readily handle interactions between the constituent ICL features. Secondly, with enough aggregations of a suitable type, lossless compression is possible to enable interactions, at least in principle. For continuous aggregators (Corso et al., 2020), the required cardinality is equal to the number of real-valued elements being aggregated (not surprising). But if arbitrary discontinuities are allowed (which is admittedly not practically feasible to implement), even just a single aggregation is technically capable of reconstructing the full set (Rubio-Madrigal & Burkholz, 2026); this conclusion follows directly from the Kolmogorov-Arnold representation theorem (Schmidt-Hieber, 2021).

Interaction experiments. We conduct multiple tests to explore the impact of interaction terms; see Appendix G for full setup details. For the first, we generalize the setup of (7) to include different interaction types involving a_i and b_i; we then compare RDBLearn against the top non-TFM-based RDB FMs, KumoRFM-2 and RT-J. Figure 3(*left*) shows the prediction results as the width of $\boldsymbol{T}^2$ is increased by adding uninformative columns that raise estimation difficulty (i.e., models are not given which interactions are useful). RT-J does not work well, possibly because the implicit graph is disconnected (see Section 5.3). Meanwhile, RDBLearn is consistently competitive with KumoRFM-2, despite g_{pf} not explicitly containing most interaction terms; see Appendix G for further analysis. For the second test, we begin with the rel-trial dataset (Robinson et al., 2024) and compute new GT labels via interactions between auxiliary table columns. Results are shown in Figure 3(*right*), where RDBLearn again outperforms the others (agentic results on this subplot are addressed in Section 7).

[3]This is admittedly difficult to quantify, as we usually only have reported results on final predictive accuracy.

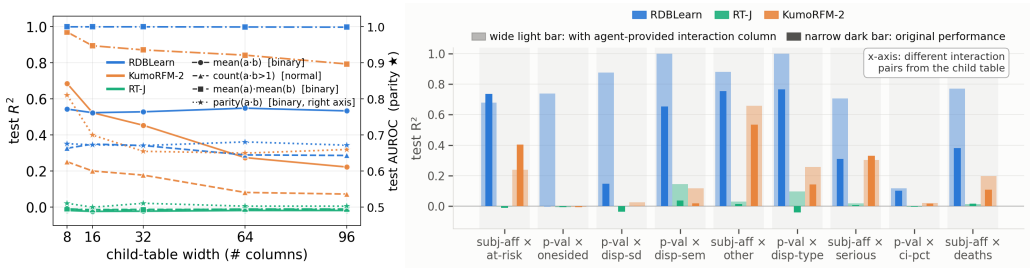


Figure 3: Prediction quality over synthetic interaction data (*left*) and quasi-real world data (*right*).

7 TFMs Form Natural Partnering with Agentic Systems

Feature discrimination in three stages. RDB predictive modeling depends heavily on the ability to differentiate, on a task-specific basis, informative versus non-informative features, relational or otherwise. There are (at least) three candidate stages where such feature screening may plausibly take place: symbolically prior to any FM application, via *reasoning agents* with access to task-dependent meta-data or domain expertise (*stage 1*); locally at the RDB encoder level using whatever context is available within $\mathcal{G}_H(\boldsymbol{x})$ (*stage 2*); and globally at the decoder level in a purely data-driven manner, using the full range of labeled ICL samples (*stage 3*). It is our assertion that TFM-based models are particularly adept at utilizing all three stages as *complements*, particularly because their g_{pf} can reduce potential interference between stages. In support of this conjecture, we offer the following:

Stage 2 as a conduit from the symbolic to the data-driven. Given the previously stated capabilities of stages 1 and 3, an ideal role for stage 2 is relatively transparent: Make sure any features that stage 1 judges to have potential are organized, to the extent possible, so as to align with the distribution of single-table features for which the subsequent stage 3 ICL decoder has been optimized. Among other things, this entails that if a particular set of features are designated as possibly useful by stage 1, they are not mixed together; some may be valuable, others ultimately not, but that is for stage 3 to decide in a data-driven manner (per results from Xu et al. (2026), this task is easier when column-wise axis-alignment is maintained). Most notably, a so-called "golden feature" with designated predictive value by a stage 1 agent should be passed directly through stage 2 without alteration or attenuation. Later we present an empirical example involving interaction terms surfaced by stage 1.

TFMs innately honor stage 1–3 roles. When we adopt TFMs, equipped with an appropriate g_{pf}, any designated golden feature from stage 1 can simply be appended to others in stage 2 without deviating from the original tabular setting under which stage 3 pre-training was conducted. But this need not be the case for existing RDB FMs with parameterized g_ϕ. For the latter, there is no guarantee that *any* input dimension/feature, golden or not, will be preserved within the output embedding. Moreover, if we simply append such features to the encoder output (bypassing the parameterized encoder itself), we then introduce potential mismatch from the original RDB-specific pre-training regime of these architectures. There is also the larger related issue of *interpretability*. An agent can directly reason over the SQL-based aggregations used to form g_{pf} and offer suggestions regarding what useful factors might be missing. In contrast g_ϕ is much more opaque since the output embeddings result from stacked Transformer layers, e.g., there is no obvious way for an agent to tell if a valuable interaction term is attenuated or absent altogether.

Empirical corroboration. Consider the hypothetical scenario whereby an agent with domain knowledge materializes a golden feature; namely, an interaction term which determines the downstream label is added as a new column. Figure 3(*right*) results show that RDBLearn leverages this information to consistency boost performance, more so than other frameworks, in line with the three-stage perspective from above. We further explore real-world implications in Appendix H, whereby RDBLearn with a single coding-agent prompt can match the performance of a SOTA multi-agent system (Huang et al., 2026) that trains numerous models *supervised* on a per-task basis.

8 Take-Home Message

Despite never seeing relational data (real or synthetic) during their pre-training stage, empirical results indicate that TFMs paired with a simple g_{pf} are competitive with the best alternative RDB FM architectures. Although our core analyses provide multiple complementary rationales for this phenomenon, we readily concede that advancements in the field may plausibly shift the balance moving forward, hopefully catalyzed in part by our findings herein.

AI Use Statement

We did *not* use AI to conceive of the ideas in this work, nor to write any of the text or technical derivations (with the exception of AI use for some last-minute polishing of a couple appendix subsections). For experiments in Sections 4 and 6, AI is used for reconstructing datasets from existing benchmarks. For experiments in Section 7, AI is used to generate SQL queries. Outside of experiments, AI was used for checking/confirming proofs (supplying minor corrections), providing a high-level paper review (e.g., correcting grammar, re-phrasing several sentences, nothing more), and producing figures and tables based on provided data.

Ethics Statement

The central goal of our paper is to increase understanding of a class of foundation models applied to relational data, and generally speaking, reducing the uncertainty of models in common use is a good thing. Overall, there is nothing in particular that stands out as having significantly broader ethical implications.

Reproducibility Statement

Beyond high-level descriptions in the main text, we have included a number of appendices describing the setup of our experiments. In particular, full details behind Figures 1–3, as well as the agentic testing in Section 7 can be found in Appendices E, F, G, and H, respectively. We have also provided an anonymous link to source code at the beginning of Appendix E. Meanwhile, all technical proofs are deferred to Appendix I

References

Kazi F Akhter, Bharath Ajendla, and Manar D Samad. A fair benchmarking of deep relational database learning models. *arXiv preprint arXiv:2607.03659*, 2026.

Tianlang Chen, Charilaos Kanatsoulis, and Jure Leskovec. RelGNN: Composite message passing for relational deep learning. *arXiv preprint arXiv:2502.06784*, 2025.

Zhikai Chen, Junyu Yin, Jialiang Gu, Siheng Xiong, Xiaoze Liu, Ruowang Zhang, Keren Zhou, and Kai Guo. OpenRFM: Dissecting relational in-context learning. *arXiv preprint arXiv:2606.04320*, 2026.

Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33:13260–13271, 2020.

Milan Cvitkovic. Supervised learning on relational databases with graph neural networks. *arXiv preprint arXiv:2002.02046*, 2020.

Amit Daniely and Eran Malach. Learning parities with neural networks. *Advances in Neural Information Processing Systems*, 33:20356–20365, 2020.

DM Dave, B Todd, and Will Cukierski. Acquire valued shoppers challenge, 2014. URL `https://kaggle.com/competitions/acquire-valued-shoppers-challenge`.

Vijay Prakash Dwivedi, Sri Jaladi, Yangyi Shen, Federico López, Charilaos I Kanatsoulis, Rishi Puri, Matthias Fey, and Jure Leskovec. Relational graph transformer. *arXiv preprint arXiv:2505.10960*, 2025.

Moonjung Eo, Min-Kook Suh, Hye-Seung Cho, Jiwon Kim, Seoyoon Kim, Sangjun Nam, and Soonyoung Lee. EXAONE Tabular 1.0: Technical Report. *arXiv preprint arXiv:2608.25774*, 2026.

Nick Erickson, Lennart Purucker, Andrej Tschalzev, David Holzmüller, Prateek Desai, David Salinas, and Frank Hutter. TabArena: A living benchmark for machine learning on tabular data. *Advances in Neural Information Processing Systems*, 38, 2026.

Matthias Fey, Weihua Hu, Kexin Huang, Jan Eric Lenssen, Rishabh Ranjan, Joshua Robinson, Rex Ying, Jiaxuan You, and Jure Leskovec. Relational deep learning: Graph representation learning on relational databases. *arXiv preprint arXiv:2312.04615*, 2023.

Matthias Fey, Vid Kocijan, Federico Lopez, J Lenssen, and Jure Leskovec. KumoRFM: A foundation model for in-context learning on relational data. *Online Tech Report*, 2025.

Hector Garcia-Molina, Jeffrey D. Ullman, and Jennifer Widom. *Database Systems: The Complete Book*. Prentice Hall, 2009.

Léo Grinsztajn, Klemens Flöge, Oscar Key, Felix Birkel, Philipp Jund, Brendan Roof, Mihir Manium, Shi Bin Hoo, Magnus Bühler, Anurag Garg, Dominik Safaric, Jake Robertson, Benjamin Jäger, Simone Alessi, Adrian Hayler, Vladyslav Moroshan, Lennart Purucker, Philipp Singer, Alan Arazi, Julien Siems, Jan Hendrik Metzen, Georg Grab, Nick Erickson, Siyuan Guo, Eliott Kalfon, Simon Bing, David Salinas, Clara Cornu, Lilly Charlotte Wehrhahn, Diana Kriuchkova, Kursat Kaya, Lydia Sidhoum, Marie Salmon, Jerry Chen, Madelon Hulsebos, Yann LeCun, Samuel Müller, Bernhard Schölkopf, Sauraj Gambhir, Noah Hollmann, and Frank Hutter. TabPFN-3: Technical report. *arXiv preprint arXiv:2605.13986*, 2026.

Prince Grover, Julia Xu, Justin Tittelfitz, Anqi Cheng, Zheng Li, Jakub Zablocki, Jianbo Liu, and Hao Zhou. Fraud dataset benchmark and applications. *arXiv preprint arXiv:2208.14417*, 2022.

Justin Gu, Rishabh Ranjan, Charilaos Kanatsoulis, Haiming Tang, Martin Jurkovic, Valter Hudovernik, Mark Znidar, Pranshu Chaturvedi, Parth Shroff, Fengyu Li, et al. RelBench v2: A large-scale benchmark and repository for relational data. *arXiv preprint arXiv:2602.12606*, 2026.

Sungwon Han, Jinsung Yoon, Sercan O Arik, and Tomas Pfister. Large language models can automatically engineer features for few-shot tabular learning. *arXiv preprint arXiv:2404.09491*, 2024.

Adrian Hayler, Klemens Flöge, Alan Arazi, Rishabh Ranjan, Jure Leskovec, Felix Birkel, Brendan Roof, Anurag Garg, Kristina Collins, Lydia Sidhoum, et al. Advancing open and reproducible relational learning: RelArena-α, TabPFN-Rel and RPI. *arXiv preprint arXiv:2608.16319*, 2026.

Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A Transformer that solves small tabular classification problems in a second. *arXiv preprint arXiv:2207.01848*, 2022.

Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeister, and Frank Hutter. Accurate predictions on small data with a tabular foundation model. *Nature*, 637(8045):319–326, 2025.

Xingyue Huang, Louis Tichelman, Jinwoo Kim, Krzysztof Olejniczak, and İsmail İlkan Ceylan. RelAgent: LLM agents as data scientists for relational learning. *arXiv preprint arXiv:2605.07840*, 2026.

Valter Hudovernik, Federico López, Vid Kocijan, Akihiro Nitta, Jan Eric Lenssen, Jure Leskovec, and Matthias Fey. KumoRFM-2: Scaling foundation models for relational learning. *arXiv preprint arXiv:2604.12596*, 2026.

Benjamin Jäger, Nick Erickson, Léo Grinsztajn, Felix Birkel, Klemens Flöge, Oscar Key, Kürşat Kaya, Jonas Kübler, Adèle Frankel, Tobias Schröder, Anurag Garg, Jan Hendrik Metzen, David Salinas, Simon Bing, Kristina Collins, Tuana Çelik, Vahid Balazadeh, Lydia Sidhoum, Tomás Pereda, Brendan Roof, Andrej Tschalzev, Siyuan Guo, Philipp Singer, Lennart Purucker, Jake Robertson, Marie Salmon, Philipp Jund, Jerry Chen, Diana Kriuchkova, Arthur Cahu, Eliott Kalfon, Adrian Hayler, Georg Grab, Vitor Monteiro, Lilly Wehrhahn, Dominik Safaric, Clara Cornu, Alan Arazi, Rylee Grace, Simone Alessi, Mihir Manium, Bernhard Schölkopf, Yann LeCun, Madelon Hulsebos, Sauraj Gambhir, Noah Hollmann, and Frank Hutter. TabPFN-3.5: Technical report. *arXiv preprint arXiv:2609.17895*, 2026.

QU Jingang, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. TabICL: A tabular foundation model for in-context learning on large data. In *International Conference on Machine Learning*, 2025.

Kyungho Kim, Geon Lee, Juyeon Kim, Dongwon Choi, Shinhwan Kang, and Kijung Shin. ReFuGe: Feature generation for prediction tasks on relational databases with LLM agents. In *Proceedings of the ACM Web Conference*, pp. 8501–8504, 2026.

Tassilo Klein, Clemens Biehl, Margarida Costa, Andre Sres, Jonas Kolk, and Johannes Hoffart. SALT: Sales autocompletion linked business tables dataset. In *NeurIPS 2024 Third Table Representation Learning Workshop*, 2024.

Vignesh Kothapalli, Rishabh Ranjan, Valter Hudovernik, Vijay Prakash Dwivedi, Johannes Hoffart, Carlos Guestrin, and Jure Leskovec. PluRel: Synthetic data unlocks scaling laws for relational foundation models. *arXiv preprint arXiv:2602.04029*, 2026.

Jan Motl and Oliver Schulte. The CTU Prague relational learning repository. *arXiv preprint arXiv:1511.03086*, 2015.

Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Transformers can do Bayesian inference. *arXiv preprint arXiv:2112.10510*, 2021.

Thomas Nagler. Statistical foundations of prior-data fitted networks. In *International Conference on Machine Learning*, pp. 25660–25676. PMLR, 2023.

Jianmo Ni, Jiacheng Li, and Julian McAuley. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pp. 188–197, 2019.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, Cambridge, MA, USA, 2017.

Lennart Purucker, Andrej Tschalzev, Nick Erickson, Gioia Blayer, David HolzmÃ1/4ller, Alan Arazi, Alexander Pfefferle, Mustafa Tajjar, GaÃ Varoquaux, Frank Hutter, et al. Beyond iid: How general are tabular foundation models, really? *arXiv preprint arXiv:2606.30410*, 2026.

Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. TabICLv2: A better, faster, scalable, and open tabular foundation model. *arXiv preprint arXiv:2602.11139*, 2026.

Rishabh Ranjan, Valter Hudovernik, Mark Znidar, Charilaos Kanatsoulis, Roshan Upendra, Mahmoud Mohammadi, Joe Meyer, Tom Palczewski, Carlos Guestrin, and Jure Leskovec. Relational Transformer: Toward zero-shot foundation models for relational data. *International Conference on Learning Representations*, 2026a.

Rishabh Ranjan, Vignesh Kothapalli, Harshvardhan Agarwal, Charilaos I Kanatsoulis, Roshan Reddy Upendra, Tom Palczewski, Carlos Guestrin, and Jure Leskovec. Large-scale pretraining unlocks few-shot prediction for relational data. In *Workshop on Graph Foundation Models: A New Era for Graph Machine Learning*, 2026b.

Joshua Robinson, Rishabh Ranjan, Weihua Hu, Kexin Huang, Jiaqi Han, Alejandro Dobles, Matthias Fey, Jan Eric Lenssen, Yiwen Yuan, Zecheng Zhang, et al. RelBench: A benchmark for deep learning on relational databases. *Advances in Neural Information Processing Systems*, 37:21330–21341, 2024.

Celia Rubio-Madrigal and Rebekka Burkholz. Fixed aggregation features can rival gnns. *arXiv preprint arXiv:2601.19449*, 2026.

Johannes Schmidt-Hieber. The Kolmogorov–Arnold representation theorem revisited. *Neural networks*, 137:119–126, 2021.

Yefan Tao, Xiyuan Zhang, Xinyi Liu, Boran Han, Danielle Maddix, Haoyang Fang, Zhen Han, Jiading Gai, Xuanqing Liu, Michael Bohlke-Schneider, et al. Mitra-v2 technical report. *arXiv preprint arXiv:2609.04540*, 2026.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 Technical Report. *arXiv preprint arXiv:2503.19786*, 2025.

Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014.

Minjie Wang, Quan Gan, David Wipf, Zhenkun Cai, Ning Li, Jianheng Tang, Yanlin Zhang, Zizhao Zhang, Zunyao Mao, Yakun Song, Yanbo Wang, Jiahang Li, Han Zhang, Guang Yang, Xiao Qin, Chuan Lei, Muhan Zhang, Weinan Zhang, Christos Faloutsos, and Zheng Zhang. 4DBInfer: A 4D benchmarking toolbox for graph-centric predictive modeling on RDBs. *Advances in Neural Information Processing Systems*, 37:27236–27273, 2024.

Penghui Wang, Wei Liu, Hong Wang, Chengyue Huang, Yuxi Sun, Zirui Wang, Hongming Huang, Quan Wang, Chunxiao Liu, et al. Xiaomi-TabLDM: A tabular foundation model technical report. *arXiv preprint arXiv:2609.03880*, 2026a.

Yanbo Wang, Xiyuan Wang, Quan Gan, Minjie Wang, Qibin Yang, David Wipf, and Muhan Zhang. Griffin: Towards a graph-centric relational database foundation model. In *International Conference on Machine Learning*, 2025.

Yanbo Wang, Jiaxuan You, Chuan Shi, and Muhan Zhang. Relational in-context learning via synthetic pre-training with structural prior. In *International Conference on Machine Learning*, 2026b.

Yangkun Wang, Jiarui Jin, Weinan Zhang, Yongyi Yang, Jiuhai Chen, Quan Gan, Yong Yu, Zheng Zhang, Zengfeng Huang, and David Wipf. Why propagate alone? Parallel use of labels and features on graphs. *International Conference on Learning Representations*, 2022.

Fang Wu, Vijay Prakash Dwivedi, and Jure Leskovec. Large language models are good relational learners. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025.

Marek Wydmuch, Łukasz Borchmann, and Filip Graliński. Tackling prediction tasks in relational databases with LLMs. *arXiv preprint arXiv:2411.11829*, 2024.

Linjie Xu, Yanlin Zhang, Quan Gan, Minjie Wang, and David Wipf. No need to train your RDB foundation model. *International Conference on Machine Learning*, 2026.

Han Zhang, Quan Gan, David Wipf, and Weinan Zhang. GFS: Graph-based feature synthesis for prediction over relational databases. *arXiv preprint arXiv:2312.02037*, 2023a.

Xingxuan Zhang, Gang Ren, Han Yu, Hao Yuan, Hui Wang, Jiansheng Li, Jiayun Wu, Lang Mo, Li Mao, Mingchao Hao, et al. Limix: Unleashing structured-data modeling capability for generalist intelligence. *arXiv preprint arXiv:2509.03505*, 2025a.

Xingxuan Zhang, Gang Ren, Hao Yuan, Hao Zou, Hongze Tan, Hui Wang, Jianhao Song, Jiansheng Li, Jiayao Zhang, Jinghan Zhang, et al. LimiX-2: A contextual mechanism network towards general structured-data intelligence. *arXiv preprint arXiv:2609.17488*, 2026.

Xiyuan Zhang, Danielle C Maddix, Junming Yin, Nick Erickson, Abdul Fatir Ansari, Boran Han, Shuai Zhang, Leman Akoglu, Christos Faloutsos, Michael W Mahoney, et al. Mitra: Mixed synthetic priors for enhancing tabular foundation models. *arXiv preprint arXiv:2510.21204*, 2025b.

Yanlin Zhang, Ning Li, Quan Gan, Weinan Zhang, David Wipf, and Minjie Wang. Elf-Gym: Evaluating large language models generated features for tabular prediction. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 5420–5424, 2024.

Zizhao Zhang, Yi Yang, Lutong Zou, He Wen, Tao Feng, and Jiaxuan You. Rdbench: ML benchmark for relational databases. *arXiv preprint arXiv:2310.16837*, 2023b.

Roman Zykov, Noskov Artem, and Anokhin Alexander. Retailrocket recommender system dataset, 2022. URL `https://www.kaggle.com/dsv/4471234`.

A Further Analysis of the Continuum Perspective from Section 4

In this section we demonstrate that simple ensembling or interpolation strategies, designed to combine canonical TFMs (without g_{pf}) and RDB-native FMs, may struggle to handle data along a continuum between single-tables and RDBs. Later, we discuss further benchmarking implications of the continuum perspective.

A.1 Ensembling Different FM Types Cannot Resolve the Continuum

One candidate workaround for navigating varying relational effects on a continuum would be to invoke both a typical RDB FM and a single-table FM, and then simply use the validation set to either pick the best one or form some type of ensemble. To analyze this possibility, we define predictions from tabular and RDB FMs respectively as

$$\begin{aligned} \hat{y}_{\text{test}}^{\text{tab}} &:= f_\theta^{\text{tab}}\big(\{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big) \;\text{ s.t. } \boldsymbol{z} = \boldsymbol{x} \\ \hat{y}_{\text{test}}^{\text{rdb}} &:= f_\theta^{\text{rdb}}\big(\{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big) \;\text{ s.t. } \boldsymbol{z} = g_\phi\big(\mathcal{G}_H(\boldsymbol{x})\big) \end{aligned} \tag{8}$$

for an arbitrary input RDB. The TFM from (8) only has access to $\boldsymbol{T}^1$ information, or any tables that can be trivially joined to $\boldsymbol{T}^1$ without a true encoder g_{pf} (i.e., by design we are assuming standard single-table operational mode here, not a TFM equipped with g_{pf} for handling rich relational effects). To formalize things further, we require precise specifications of a suitable RDB generation process. We must also introduce some form of practical constraint on the RDB FM encoder; otherwise, it is trivially a superset of the TFM. We present these definitions first, before proceeding to the main result of this section.

Definition A.1. We define $\mathcal{D}_\tau$ as an RDB induced by a single table via the following generation process p_τ, where continuation paramter $\tau \in [0, 1]$ mitigates the significance of cross-table relational factors. First, we draw tabular rows $\boldsymbol{x}_i = [\boldsymbol{x}_i^{\text{real}}; \boldsymbol{x}_i^{\text{cat}}]$ from some $p(\boldsymbol{x}_i)$, where $\boldsymbol{x}_i^{\text{real}} \in \mathbb{R}^d$ represent numerical features, $\boldsymbol{x}_i^{\text{cat}} \in \mathcal{C}^\kappa$ are categorical features with $|\mathcal{C}| \leq B$ and $i = 0, \ldots, n$. We assume that the marginal $p\left(\boldsymbol{x}^{\text{real}}\right)$ is an absolutely continuous distribution satisfying $p\left(\boldsymbol{x}^{\text{real}}\right) \leq p_{\max}$ for all $\boldsymbol{x}^{\text{real}} \in \mathcal{R}^d$. We treat each categorical column as an FK pointing to a new auxiliary table $\boldsymbol{T}^k$ (with $k \in \{1, \ldots, \kappa\}$) composed of a single PK column (the elements of which correspond with the unique categories in the referring FK column). And finally, we generate labels by first uniformly sampling $d' < d$ indices from $\{1, \ldots, d\}$ to form the reduced numerical features $\widetilde{\boldsymbol{x}}_i^{\text{real}} \in \mathbb{R}^{d'}$. We then compute labels as

$$y_i \;=\; \psi\big[(1-\tau) f_1(\widetilde{\boldsymbol{x}}_i^{\text{real}}) + \tau f_2\big(\mathcal{G}_H(\boldsymbol{x}_i) \backslash \boldsymbol{x}_i\big)\big] \;\in\; \mathcal{Y}, \;\; \forall 0 = 1, \ldots, n, \tag{9}$$

where f_1 and f_2 are continuous functions mapping their respective domains to a shared r-dimensional embedding space $\mathbb{R}^r$, after which $\psi : \mathbb{R}^r \to \mathcal{Y}$ converts to the label space.[4] We reserve $\{\mathcal{G}_H(\boldsymbol{x}_i), y_i\}_{i=1}^n$ for ICL while treating sample $n = 0$ as $\{\mathcal{G}_H\left(\boldsymbol{x}_{\text{test}}\right), y_{\text{test}}\}$.

Definition A.2. We define a space of minimally constrained RDB encoders $\bar{g}$, elements of which produce fixed-length embeddings as $\boldsymbol{z} = g_\phi\big(\mathcal{G}_H(\boldsymbol{x})\big) \in \mathbb{R}^{d_z}$. Furthermore, for any $g_\phi \in \bar{g}$, we impose that

$$\min_\eta \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau}\left[\left\|\boldsymbol{x}_{\text{test}}^{\text{real}} - \eta\big(\{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right\|\right] > 0, \tag{10}$$

with expectation over the generation process from Definition A.1 and $\eta : \mathbb{R}^{d_z \times (n+1)} \times \mathcal{Y}^n \to d$ as an aribitrary continuous mapping.

The intuitive implication of Definition A.2 is simply that $\boldsymbol{x}_{\text{test}}^{\text{real}}$, which may be high-dimensional, cannot be completely reconstructed (in expectation) using all of the encoded information available to an ICL-based decoder q_θ, namely (namely $\{\boldsymbol{z}_i, y_i\}_{i=1}^n$ and $\boldsymbol{z}_{\text{test}}$). This is nominally a reasonable restriction to impose given that typical RDB encoders must ultimately compress into fixed-length embeddings, and such compression may apply to an original feature row $\boldsymbol{x}$, particularly so for any high-dimensional numerical values. Moreover, (10) does *not* rule out complete reconstruction of

[4] For reference, $\mathcal{G}_H(\boldsymbol{x}) \backslash \boldsymbol{x}$ represents all H-hop neighborhood information available at the timestamp associated with $\boldsymbol{x}$ (following the newly-generated PK-FK pairs), but excluding the seed row features $\boldsymbol{x}$.

$d-1$ elements of $\boldsymbol{x}_{\text{test}}^{\text{real}}$, indicating minimal enforced compression. Given these definitions, we now analyze the limitation of ensembling:

Proposition A.3. *For any $\varepsilon \in (0,1)$ and $C > 0$ there will always exist fixed mappings $\{\Psi', f_1', f_2'\}$ such that the following simultaneously hold:*

$$\min_{f_\theta^{tab}} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau}\Big[\, \big|y_{test} - \hat{y}_{test}^{tab}\big| \,\Big] = 0 \qquad \tau = 0$$

$$\min_{f_\theta^{rdb},\, g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau}\Big[\, \big|y_{test} - \hat{y}_{test}^{rdb}\big| \,\Big] = 0 \qquad \tau = 1 \tag{11}$$

$$\min_{f_\theta^{tab},\, f_\theta^{rdb},\, g_\phi \in \bar{g},\, \alpha \in [0,1]} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau}\Big[\, \big|y_{test} - \big(\alpha\, \hat{y}_{test}^{tab} + (1-\alpha)\hat{y}_{test}^{rdb}\big)\big| \,\Big] \geq C \qquad \forall\, \tau \in [\varepsilon, 1-\varepsilon],$$

where $\mathcal{D}_\tau \sim p_\tau$ indicates an RDB sampled according to Definition A.1 with $\{\Psi, f_1, f_2\} = \{\Psi', f_1', f_2'\}$ and $n > d'$, and $\bar{g}$ denotes any RDB encoder satisfying the conditions of Definition A.2.

Even though we are optimizing separately for each τ, this relatively straightforward result shows that any additive combination of the two FMs (as produced by typical ensembling) can perform arbitrarily badly. This is despite the fact that for the purely tabular regime ($\tau = 0$), the TFM exhibits ideal behavior, while for the purely relational case (i.e., $\tau = 1$ implying no dependence on the original target row features) the RDB FM performs perfectly.

Admittedly, we can always try training a more complex combining function to address this scenario, but then we are diverging from the targeted FM regime to supervised learning. In contrast, if we pre-train our original FMs to consistently handle such cases, then in principle we should arrive at a model the operates effectively on both single and multiple tables. *But of course this is the main point we are advocating for in the first place; namely, RDB FMs should work on single tables or single-table adjacent regimes as well.*

A.2 Benchmarking Implications of the Continuum Perspective

If there were always a clear-cut distinction between single tables and RDBs, then benchmarking would be straightforward: simply partition into separate domains and evaluate accordingly, RDB FMs tested only on RDBs (e.g., RelArena, RelBench, 4DBInfer, etc.) and TFMs only on single tables (e.g., TabArena). But in light of the continuum described in Section 4, is such strict partitioning unequivocally warranted?

For example, consider breaking out of this mold by applying an RDB FM to a suite of traditional tabular benchmarks. If performance is competitive relative to frontier TFMs that specialize on single tables, then our confidence in the robustness of this RDB FM justifiable grows. But now suppose instead that performance is notably lagging. It is tempting to simply discount such cases as the magisteria of single-table FMs and move on; however, the continuum perspective complicates any such immediate dismissal. After all, if an RDB FM does not work well on single-tables, what about performance on RDBs along a continuum that are minimally different from single tables? More broadly, it is reasonable to question how large is the multi-table relational neighborhood around any single-table benchmark within which unsatisfactory predictions are likely to be produced. Or conversely, we may ask how large along any continuum must this neighborhood grow before the RDB FM capabilities begin to outshine TFM alternatives by exploiting relational effects. It is precisely these factors that (at least partially) motivate RDB FMs based on TFMs that already perform well on single tables.

B Implications of Encoder Label Usage

Two counter-points challenge the efficacy of the parameter-free encoders g_{pf} upon which TFMs applied to RDBs are based. Both involve the use of observable labels as inputs to a trainable g_ϕ, and the potential benefits of doing. We analyze these scenarios next, providing new insights that further inform foundational questions of RDB FM development.

B.1 Labels as Extra Discriminative Features

Per previously-stated definitions, an RDB neighborhood subgraph $\mathcal{G}_H(\boldsymbol{x})$ may include nodes associated with other rows of $\boldsymbol{T}^1 \equiv \boldsymbol{X}$, provided these rows are reachable within H PK-FK hops and have an earlier time-stamp to avoid temporal leakage. As these rows have associated labels, it is natural to consider including them as extra node features to improve downstream predictability. We adopt $\mathcal{G}^*_H(\boldsymbol{x})$ to reference the resulting subgraph with these labels concatenated as additional node features; we are assuming here that any labels included within $\mathcal{G}^*_H(\boldsymbol{x})$ have time-stamps earlier than $\boldsymbol{x}$. In more traditional supervised learning regimes involving graphs, this is a common practice for improving node classification accuracy (Wang et al., 2022) (akin to a trainable form of label propagation). More recently, analogous ideas have been applied to RDB foundation models, whereby trainable encoders are explicitly designed to process past labels of similar entities within an RDB neighborhood (Chen et al., 2026; Hudovernik et al., 2026; Ranjan et al., 2026a).

But to what degree can we actually rely on such neighborhood labels as predictive features, specifically in the context of trainable RDB foundation model encoders g_ϕ applied to $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ evaluated at an unlabeled test row of interest $\boldsymbol{x}_{\text{test}}$. To begin, for an arbitrary RDB we define $\mathcal{G}^*_{\text{test}}$ as an augmented version of $\mathcal{G}$, with $\boldsymbol{x}_{\text{test}}$ appended to $\boldsymbol{X}$, known labels associated with $\boldsymbol{X}$ included as candidate node features, but restricted to entities with timestamps earlier than $\boldsymbol{x}_{\text{test}}$. We also let $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$ denote $\mathcal{G}^*_{\text{test}} \backslash \mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$, meaning all RDB context outside of the label-infused H-hop ego-network centered at test point $\boldsymbol{x}_{\text{test}}$. And lastly, we adopt $y_H(\boldsymbol{x}_{\text{test}})$ to indicate the set of all past labels present within $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$; this of course does *not* include the unknown label y_{test} associated with $\boldsymbol{x}_{\text{test}}$ itself. We then have the following:

Proposition B.1. *Given any ego-network $\mathcal{G}^*_H(\boldsymbol{x}_{test})$ with $|y_H(\boldsymbol{x}_{test})| > 0$, distributions over remaining RDB content $p\big(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{test}), y_{test} | \mathcal{G}^*_H(\boldsymbol{x}_{test})\big)$ will always exist such that:*

1. (Foundation model setting) *y_{test} is a deterministic function of $\mathcal{G}^*_{test}$, and yet any possible fixed encoder g_ϕ and prediction head q_θ perform no better than chance estimating y_{test} in the sense that*

$$\mathbb{E}\left[\mathbb{KL}\Big[p(y_{test}|\mathcal{G}^*_{test}) \;||\; q_\theta\big(y_{test} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{test})\big]\big)\Big]\right] \;\geq\; \mathbb{E}\left[\mathbb{KL}\Big[p(y_{test}|\mathcal{G}^*_{test}) \;||\; \text{Unif}\big(y_{test}\big)\Big]\right], \tag{12}$$

   with outer expectations over $p\big(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{test})|\mathcal{G}^*_H(\boldsymbol{x}_{test})\big)$.

2. (Supervised setting) *Given $\{\mathcal{G}^*_H(\boldsymbol{x}_i), y_i\}_{i=1}^n$ and y_{test} extracted from any single draw from $p\big(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{test}), y_{test}|\mathcal{G}^*_H(\boldsymbol{x}_{test})\big)$, there exists a draw-specific $f_\theta = q_\theta \circ g$ satisfying $y_i = f_\theta\big[\mathcal{G}^*_H(\boldsymbol{x}_i)\big]$ for all i and $y_{test} = f_\theta\big[\mathcal{G}^*_H(\boldsymbol{x}_{test})\big]$, where g is a parameter-free mapping from any $\mathcal{G}^*_H(\boldsymbol{x}_{test})$ to $\mathbb{R}^m$ and q_θ is a ReLU network with finite width and depth.*

Intuition behind of Proposition B.1. Neighborhood labels within $\mathcal{G}^*_H(\boldsymbol{x})$ can be highly discriminative in a *supervised* setting, where training specifically on samples $\{\mathcal{G}^*_H(\boldsymbol{x}_i), y_i\}_{i=1}^n$ extracted from a *single* RDB predictive task facilitate learning a task-specific discriminative function f_θ applicable to test points $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ (part 2). And yet it is not generally possible to determine a stable predictive role for these labels within an encoder that must remain *fixed* across varying datasets/tasks facing a foundation model architecture of the form $q_\theta\big(y \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big)$ (part 1). Given this implication of Proposition B.1, why then are existing schema-agnostic foundation models like the original relational Transformer (RT) model from Ranjan et al. (2026a) capable of making successful predictions? A simple plausible hypothesis is that, within the handful of real-world RDBs available for pre-training RT models, homophily relationships dominate, so simply adopting the majority label of local neighbors can be a highly discriminative feature. Empirical results from Ranjan et al. (2026a)[Table 1] support this conclusion, namely, the similar performance of a basic entity mean predictor explicitly predicated on homophily (66.7 avg AUROC) relative to RT (69.7 avg AUROC).

Implications for ICL-based models. ICL provides one pathway for diversifying label usage effectively beyond the homophily-dominant regimes mentioned above; however, this remains true whether or not a parameter-free encoder g_{pf} is adopted. This is because even without trainable parameters, encoder representations can include diverse label aggregations that are sensitive to both homophily and heterophily relationships; the subsequent ICL-prediction head then sorts which are relevant for any given task. Note that canonical aggregators can trivially handle the parity check

example constructed within the proof of Proposition B.1 (see Appendix I). Hence given evidence available thus far, the use of labels as discriminative features does not by itself provide a compelling reason to favor trainable parameterized RDB encoders g_ϕ within ICL-based models. And so TFMs equipped with g_{pf} ultimately remain viable independently of whether or not such labels are available.

B.2 Labels for Adjudicating Feature Importance

Analysis from Xu et al. (2026) indicates that uninformative feature columns within RDBs cannot generally be filtered out using a dataset-agnostic encoder, which reduces the motivation for including trainable encoder parameters. However, this calculus may change if the encoder is provided with local neighborhood labels (as in $\mathcal{G}_H^*(\boldsymbol{x})$). The latter provide task-specific information, which may plausibly facilitate feature-column differentiation on a task-by-task basis as required by RDB foundation models. But is such differentiation generally possible even when $\mathcal{G}_H(\boldsymbol{x})$ is switched to $\mathcal{G}_H^*(\boldsymbol{x})$? The answer is negative in the following sense:

Proposition B.2. *Given any representative RDB ego-network $\mathcal{G}_H(\boldsymbol{x}_{test})$ constructed from non-empty tabular feature columns and $|y_H(\boldsymbol{x}_{test})| > 0$, there will always exist distributions over $\widetilde{\mathcal{G}}_H^*(\boldsymbol{x}_{test})$ and $y_H(\boldsymbol{x}_{test}) \cup y_{test}$ such that:*

1. *Each label in $y_H(\boldsymbol{x}_{test}) \cup y_{test}$ is a shared deterministic function of some support set S of tabular column features and independent of all remaining columns.*
2. *Encoder input labels $y_H(\boldsymbol{x}_{test})$ provide no information about the unknown generating support set S in the sense that $p\big(S|\mathcal{G}_H^*(\boldsymbol{x}_{test})\big) = p\big(S|\mathcal{G}_H(\boldsymbol{x}_{test})\big)$.*

Implications for ICL-based models. While not strictly ruling out any benefit of trainable encoders, Proposition B.2 does mute confidence in their ability to robustly screen uninformative columns through the use of local neighborhood labels as additional inputs. This plausibly suggests that, at least for RDBs with large numbers of feature columns (many of which are likely uninformative), trainable encoders may be unnecessary, especially given that parameter-free encoders are already structured to naturally handle column-wise feature screening (Xu et al., 2026). Therefore, consistent with our conclusions from Appendix B.1, TFMs with an appropriate g_{pf} need not fall behind, even when considering encoder label usage for assessing task-specific feature importance.

We also remark in closing that the RDB FM architecture described in Hudovernik et al. (2026) involves first converting each RDB table $\boldsymbol{T}^k$ into revised embeddings $\boldsymbol{U}^k$ using task-specific labels, prior to presentation to $\mathcal{G}_H$. While certainly an innovative idea, the extent to which uninformative columns can be resolved at this early stage may still have limitations (e.g., the counter-example used to prove Proposition B.2 would disallow filtering uninformative columns using the signal within a given $\boldsymbol{U}^k$ alone). However, further treatment of these possibilities lies beyond the scope of this work.

C Degeneracy Example from Random Walk Sampling

Assume data drawn according to Definition A.1, with original feature table $\boldsymbol{T}^1 = \boldsymbol{X}$ and corresponding labels $\boldsymbol{y}$, as well as induced auxiliary table $\boldsymbol{T}^2 \in \mathcal{C}^B$ composed of a single PK column. The latter is assumed to correspond with a categorical column in $\boldsymbol{T}^1$ that we treat as an FK column pointing to the PKs in $\boldsymbol{T}^2$. Granting this simplified setup, we observe that any random walk originating at an arbitrary row i of $\boldsymbol{T}^1$ will at most reach $|\mathcal{R}_i|$ other rows of $\boldsymbol{T}^1$. Here $\mathcal{R}_i$ denotes the set of rows within $\boldsymbol{T}^1$ that share the same FK value as row i. There is no mechanism for longer random walks to eventually extend across the entire RDB; each must stall within $\mathcal{R}_i$ because ultimately the underlying graph $\mathcal{G}$ extracted from the RDB is is not fully connected.

The graph underlying the experiments shown in Figure 3(*left*) exhibits an analogous phenomenon. The only difference is that the FK column switches from (implicitly) $\boldsymbol{T}^1$ to (explicitly) $\boldsymbol{T}^2$ in this example, while the corresponding $\mathcal{G}$ remains disconnected. This might help explain why the RT-J model does not happen to work well within the particular context of Figure 3(*left*) experiments. That being said, if the RT-J context budget is not filled by random-walk accessible rows, the fallback appears to be uniform sampling elsewhere. However, it is not clear how often this occurs in practice.

D Language Model Agent Usage

A growing body of evidence suggests that LLM agents can play a useful role in tabular model development (Han et al., 2024; Huang et al., 2026; Kim et al., 2026; Zhang et al., 2024). That being said, their usage in the specific context of RDB predictive modeling involves notable caveats from a benchmarking perspective. One in particular relates to potential leakage issues that may obscure the degree to which we can plausibly extrapolate benchmarking results to expected performance within new siloed industry settings. This concern stems from the fact that existing public benchmarks are scarce, largely because of privacy concerns with real-world data. And yet predictive solutions addressing the few RDB datasets that do exist, often involving feature engineering and human expertise, are readily visible to frontier LLM agents during their pre-training process or during their inference-time retrieval. Hence agent-based solutions may simply lean on problem-specific prior human knowledge to a degree that may not be possible on less-familiar or esoteric tasks within private RDB repositories. Of course this is not meant to discourage LLM-based RDB frameworks, and certainly TFMs equipped with g_{pf} (e.g., RDBLearn, TabPFN-Rel-OSS, etc.) are a natural fit for agentic enhancements, the topic of Section 7. We only advocate that some caution is warranted when interpreting results, especially when extrapolating to new domains. While similar concerns have certainly been raised with other data modalities, we would argue that the situation is more acute in RDB regimes with only limited, relatively well-known benchmarking sources that are susceptible to memorized solutions.

E Figure 1 Experiment Details and Full Results

For code behind Figure 1 results, please see the following anonymous link:

`https://anonymous.4open.science/status/iclr-submission-39479-35EC/`

E.1 Description of TFM-based Approaches

For TFM-based pipelines we adopt **RDBLearn** (Xu et al., 2026) and **TabPFN-Rel** (Hayler et al., 2026) as mentioned in the main text; we use **TabPFN-Rel-OSS** to refer to the open-source version of TabPFN-Rel. Since our goal is to specifically examine TFM-based approaches, we run both RDBLearn and TabPFN-Rel-OSS out of the box ourselves across all benchmarks,[5] with no dataset- or task-specific adjustments. However, we found that TabPFN-Rel-OSS did not run on the multi-class SALT benchmark in our environment; hence this is our only exclusion.[6] These approaches are conceptually the same, with a shared g_{pf} implementation, but differing primarily in which tabular base models are chosen and details of how ICL sample sets are constructed. For example, TabPFN-Rel-OSS includes validation set samples in the final chosen ICL set as in Hudovernik et al. (2026), and also oversamples ICL instances with more recent past timestamps; the latter is consistent with benefits of recency featurization bias explored by Huang et al. (2026). Neither RDBLearn nor TabPFN-Rel-OSS include text-column handling (text columns are simply ignored), unlike the closed-source TabPFN-Rel version (see below) which includes a proprietary module for this purpose.

We also remark that shortly before the ICLR submission deadline another crop of advanced single-table FMs were released, including TabPFN-v3.5 (Jäger et al., 2026), EXAONE (Eo et al., 2026), LimiX-2 (Zhang et al., 2026), Mitra-v2 (Tao et al., 2026), and Xiaomi-TabLDM (Wang et al., 2026a). These approaches can readily be combined with an encoder g_{pf} to further improve performance relative to alternative RDB FM formulations. However, we do not rerun all experiments using these options, since our central goal herein is not to surface the best possible TFM configuration within RDB operational settings. The latter represents a moving target, and remains somewhat orthogonal to our primary objective, which the results from Figure 1 already achieve.

On the practical side, we evaluate RDBLearn and TabPFN-Rel-OSS on a wider task set than the original papers considered, which requires executing these models ourselves. Both methods were run on a single Nvidia RTX 5090 (32G) GPU. For RDBLearn, we encountered OOM errors in some

[5] An exception is TabPFN-Rel-OSS specifically on RelBench-v1 tasks, whereby existing results were pulled directly from Hayler et al. (2026) rather than running ourselves.

[6] TabPFN-Rel-OSS presently does not handle class cardinalities that exceed the base model capability, and SALT tasks happen to exceed this limit for now.

tasks when the TFM is Limix or TabFM. These sweep cases are skipped, so RDBlearn just selects the best setup from the remaining healthy cases. For TabPFN-Rel-OSS, we use the default inference configuration.

E.2 Description of Other Baseline Models

To compare against the TFM-based approaches from above, we present evaluations across a wide range of alternative RDB predictive architectures, including both foundation models and more traditional supervised learning approaches. We emphasize though that results are not presently available for all models on all datasets. For example, it is frequent practice to only report results on RelBench-v1 datasets and tasks, or to only report results on some tasks within a given benchmark.

Closed-Source Foundation Models. We include the original industry model **KumoRFM-1** (Fey et al., 2025) as well as the more recent version **KumoRFM-2** (Hudovernik et al., 2026) that incorporates task-specific labels within the encoder and constituent table embeddings.[7] We also compare with **OpenRFM** (Chen et al., 2026); while this model is presumably intended to be open-sourced at some point, presently the full code is not available. For all of these closed-source approaches we rely on posted performance results from original authors for evaluation purposes, with one exception. We follow Grinsztajn et al. (2026), who have independently reevaluated KumoRFM-2 performance via API calls for certain datasets using author-provided benchmarking scripts.

Open-Source Foundation Model Baselines. Open-source baselines include the schema-agnostic **Griffin** (Wang et al., 2025), **RelLLM** (Wu et al., 2025), and **RT** (relational transformer) (Ranjan et al., 2026a) models, as well as an updated RT version enhanced with **PluRel** synthetic pretraining data (Kothapalli et al., 2026), denoted **RT-P**. On the open-source ICL-based side we include results with **RT-J**, also based on the RT architecture but with modifications such as random-walk sampling to collect ICL samples (as discussed in Section 5.3 of the main text) and pretraining with a large set of real-world RDBs called **The Join** (Ranjan et al., 2026b). Furthermore, we compare results from **RDB-PFN** (Wang et al., 2026b), with synthetic RDB pre-training, and **RandGNN** (Xu et al., 2026), which combines a random-weight GNN encoder with a TFM-based model. RandGNN represents a form of ablation w.r.t. RDBLearn and TabPFN-Rel-OSS, in that it explicitly mixes cross-column information in a manner at odds with the theoretical considerations upon which these models are based. And finally, for pure language model baselines, we evaluate **LLM-A** (Team et al., 2025) (as tested in Ranjan et al. (2026a)) and **LLM-B** (Wydmuch et al., 2024); both approaches are predicated on serialized RDB neighborhood representations and/or ICL samples. Although the underlying LLMs are not open-source, we still treat these approaches as open-source given that the prompts used to reproduce results are publicly available and no hidden fine-tuning steps are involved.

Task-Specific Supervised Models. While our emphasis is on RDB foundation models, it is still useful to contextualize FM performance results w.r.t. methods benefitting from per-task supervised training. On this front, largely following prior work, we include the more classical standards **LightGBM**, **XGBoost**, and **GraphSage** as evaluated in Hudovernik et al. (2026). Meanwhile, for recent SOTA supervised architectures we select **RelGNN** (Chen et al., 2025), **RelGT** (Dwivedi et al., 2025), **RelAgent** (Huang et al., 2026), and **RT-P+fine-tuning**; the latter combines PluReL-based pretraining and task-specific fine-tuning with the RT-J architectural enhancements mentioned previously to achieve SOTA performance on RelArena-α (Hayler et al., 2026).

E.3 RDB Benchmarks

In producing Figure 1, we evaluate RDB predictive modeling approaches over a diverse suite of benchmarks, including regression and binary classification tasks from both **RelBench-v1** (Robinson et al., 2024) and **RelBench-v2** (Gu et al., 2026). Note that RelBench-v1 datasets and tasks are also equivalent to the recently-posted **RelArena-α** benchmarking platform (Hayler et al., 2026) patterned after TabArena (Erickson et al., 2026). The primary difference is that RelArena-α enforces more rigid/standardized submission protocols that some published work on RelBench-v1 may not satisfy.

[7] KumoRFM-2 has generally been treated as a closed-source model and we maintain this designation here given that all KumoRFM-2 results in Figure 1 are based on private API calls. Moreover, the KumoRFM-2 pre-training regime and extent of real-world RDB exposure remain undisclosed.

That being said, RDBLearn and TabPFN-Rel-OSS already adhere to these guidelines, and it is only to the advantage of other models with which we compare if they do not. This means our results in Figure 1 are actually conservative in showcasing the abilities of TFM-based approaches to RDB predictive modeling.

Beyond the above, we test over five **4DBInfer** (Wang et al., 2024) binary classification datasets that have been used in several recent works (Huang et al., 2026; Hudovernik et al., 2026), as well as the multi-class predictive tasks from the **SALT** repository (Klein et al., 2024) as adopted by Hudovernik et al. (2026). Standard train/dev/test splits are used for all experiments. Note that for RelBench-v1, we omit only rel-f1 and rel-event datasets because of potential leakage issues mentioned in prior work, and we do not have results for all baseline models on these datasets.

E.4 Extended Empirical Results

In this section we provide the full set of experimental results upon which the final rankings shown in Figure 1 from the main text are based. Overall, the TFM-based methods (along with their attendant g_{pf}) achieve the best overall performance among RDB FMs. This follows by comparing numbers for RDBLearn, TabPFN-Rel-OSS, and TabPFN-Rel versus other RDB FMs in the tables below. Surprisingly, these frameworks even exceed the results produced by many per-task supervised models.

Table 1: Results on 4DBInfer binary classification tasks (AUC↑). Bold = best per column, underline = second best.

Method	AB churn	OB ctr	RR cvr	SE upvote	SE churn	**Avg AUROC** ↑	**Avg Rank** ↓
Supervised							
XGBoost	50.00	50.00	50.00	49.68	50.84	50.10	9.6
GNN	75.71	62.39	84.70	88.61	85.58	79.40	4.6
RelAgent	**79.44**	62.63	86.35	**89.45**	<u>89.03</u>	81.38	**2.0**
Foundation model							
Closed-Source							
KumoRFM-1	75.59	61.11	84.36	<u>88.92</u>	87.87	79.57	5.1
KumoRFM-2	77.06	61.97	84.36	88.50	87.92	79.96	4.9
RDB-PFN	73.38	52.59	78.12	85.60	85.50	75.04	7.2
Open-source							
Griffin	53.32	52.22	39.92	71.66	72.18	57.86	8.8
RandGNN	77.42	49.29	66.21	61.06	77.41	66.28	7.6
TabPFN-Rel-OSS	77.21	<u>63.68</u>	<u>88.92</u>	88.59	**89.14**	<u>81.51</u>	<u>2.6</u>
RDBLearn	<u>78.71</u>	**63.76**	**89.06**	87.86	88.69	**81.62**	<u>2.6</u>

Table 2: Results on the SALT benchmark (MRR↑). Bold = best per column; underline = second best.

	salt-item			salt-sales					Avg. MRR	Avg. Rank
Method	Plant	Shipping Point	Incoterm	Incoterm	Office	Group	Payment Terms	Shipping Condition	↑	↓
Supervised										
GraphSAGE	<u>0.99</u>	<u>0.98</u>	0.69	0.62	**1**	0.16	0.37	0.57	0.67	3.88
LightGBM	0.61	0.28	0.73	0.73	<u>0.99</u>	0.02	0.10	0.51	0.50	4.75
Foundation model										
Closed-Source										
KumoRFM-1	<u>0.99</u>	**0.99**	<u>0.79</u>	<u>0.81</u>	**1**	0.38	0.38	0.78	0.77	<u>2.56</u>
KumoRFM-2	<u>0.99</u>	<u>0.98</u>	0.78	<u>0.81</u>	**1**	<u>0.61</u>	<u>0.61</u>	<u>0.79</u>	<u>0.82</u>	<u>2.56</u>
Open-source										
RDBLearn	**1**	**0.99**	**0.82**	**0.91**	**1**	**0.69**	**0.79**	**0.84**	**0.88**	**1.25**

Table 3: Results on Relbench v2 binary classification tasks (AUROC↑). Bold = best per column; underline = second best.

	mimic	ratebeer			arxiv	**Avg AUROC**	**Avg Rank**
Method	stay	beer-churn	user-churn	dormant	citation	↑	↓
Supervised							
LightGBM	51.81	76.21	83.92	75.79	71.21	71.79	5.80
GraphSAGE	55.01	78.67	94.27	80.51	82.50	78.19	3.60
Rel-Agent†	–	84.70	97.43	**83.33**	82.61	–	–
Foundation model							
Closed-Source							
KumoRFM-1	**56.33**	75.06	91.38	77.10	80.62	76.10	4.40
KumoRFM-2	55.28	83.84	97.43	80.65	81.71	79.78	3.00
OpenRFM (TabICL)†	–	79.60	–	77.50	81.00	–	–
OpenRFM (TabPFN)†	–	80.90	–	79.10	81.00	–	–
Open-source							
TabPFN-Rel-OSS	53.10	80.15	98.28	82.32	**82.69**	79.31	2.40
RDBLearn	55.36	**85.15**	**98.48**	81.85	82.37	**80.64**	**1.80**

† Rel-Agent and OpenRFM are excluded when calculating average AUROC and rank because results are missing in some cases (Rel-Agent: no mimic-stay; OpenRFM: no mimic-stay or ratebeer user-churn). Note also that public code is not available to produce results for the missing tasks at the time of this writing.

Table 4: Results on Relbench v2 regression tasks (MAE↓). Bold = best per column, underline = second best.

Method	`ratebeer` `user-count`	`arxiv` `publication`	**Geo. Mean** ↓	**Avg. Rank** ↓
Baseline				
Global Median	15.124	0.577	2.954	7.75
Entity Median	13.079	0.874	3.381	8.00
Supervised				
LightGBM	20.350	0.577	3.427	8.25
GraphSAGE	7.374	0.513	1.945	5.00
RelAgent	**6.021**	0.462	**1.668**	**1.50**
Foundation model				
Closed-Source				
KumoRFM-1	11.063	0.518	2.394	6.00
KumoRFM-2	7.298	0.487	1.885	3.50
Open-source				
TabPFN-Rel-OSS	7.131	0.499	1.887	3.50
RDBLearn	6.870	**0.459**	1.776	**1.50**

Table 5: Results on Relbench v1 binary classification tasks (AUROC↑). Bold = best per column, underline = second best.

	avito		trial	amazon		stack		hm	**Avg AUROC**	**Avg Rank**
Method	click	visit	out	user	item	eng	badge	churn	↑	↓
Supervised										
LightGBM	53.60	53.05	70.09	52.22	62.54	63.39	63.43	55.21	59.19	19.12
GraphSAGE	65.90	66.20	68.60	70.42	82.81	90.59	88.86	69.88	75.41	6.12
RelGT	<u>68.30</u>	66.78	68.61	70.39	82.55	90.53	86.32	69.27	75.34	6.69
RelGNN	68.23	66.18	71.24	<u>70.99</u>	82.64	**90.75**	<u>88.98</u>	<u>70.93</u>	76.24	4.25
RelAgent	**68.36**	<u>67.79</u>	71.86	70.78	<u>82.84</u>	90.41	88.42	**71.07**	<u>76.44</u>	**3.50**
RT-P+fine-tuning	58.34	67.09	72.35	**71.35**	**83.27**	89.68	**89.16**	70.44	75.21	5.25
Foundation model										
Closed-Source										
KumoRFM-1	64.85	64.11	70.79	67.29	79.93	87.09	80.00	67.71	72.72	11.56
KumoRFM-2 (orig.)*	67.42	**69.41**	72.03	69.10	82.17	89.40	87.15	69.27	75.74	6.75
KumoRFM-2 (repro.)*	67.42	**69.41**	72.03	67.71	80.18	88.69	85.40	67.81	74.83	7.94
TabPFN-Rel	67.52	66.80	**76.47**	70.86	82.80	<u>90.60</u>	88.04	70.52	**76.70**	<u>3.75</u>
OpenRFM (TabICL)†	62.30	64.60	60.80	–	–	88.50	86.30	68.60	–	–
OpenRFM (TabPFN)†	63.90	64.60	63.40	–	–	88.40	86.40	68.90	–	–
Open-source										
LLM-A	59.80	62.70	57.40	58.10	62.10	78.00	80.00	59.80	64.74	17.31
LLM-B	61.32	60.28	55.72	60.56	71.96	81.01	71.13	64.34	65.79	17.00
RelLLM	62.28	56.17	59.02	60.07	64.10	69.46	62.12	55.95	61.15	18.12
Griffin	45.90	60.70	51.00	62.30	69.00	77.50	73.50	60.20	62.51	18.62
RT	59.50	61.80	51.80	64.00	70.90	75.70	80.10	62.80	65.83	16.94
RT-J	59.14	61.96	66.46	68.73	78.98	83.94	75.25	67.13	70.20	14.25
RT-P	47.90	63.40	51.80	65.00	72.50	86.20	82.00	66.00	66.85	15.56
RDBPFN	64.46	65.82	64.50	65.92	79.59	88.04	84.52	67.21	72.51	11.75
RandGNN	61.58	64.15	55.98	65.87	77.73	85.14	83.14	66.54	70.02	13.62
TabPFN-Rel-OSS	61.45	66.88	<u>73.06</u>	70.24	82.79	90.58	86.35	70.57	75.24	5.75
RDBLearn	67.55	66.71	72.71	69.33	82.25	89.92	85.23	68.27	75.25	7.12

* Results are drawn from both the original KumoRFM-2 paper (Hudovernik et al., 2026), as well as an independent reproduction from Grinsztajn et al. (2026) using KumoRFM-2 API calls for some datasets.
† OpenRFM is excluded when calculating average AUROC and rank because results are not available for all datasets, and public code is not available to produce results for the missing tasks.

Table 6: Results on Relbench-v1 regression tasks (MAE↓). Bold = best per column; underline = second best.

	avito	trial		amazon		stack	hm	**Geo. Mean**	**Avg Rank**
Method	ctr	adverse	success	user	item	votes	sales	↓	↓
Base									
Global Median	0.043	57.533	0.462	16.783	64.234	0.068	0.076	1.303	14.29
Entity Median	0.046	57.930	0.441	17.423	66.436	0.069	0.078	1.329	15.43
Supervised									
LightGBM	0.041	44.011	0.425	16.783	60.569	0.068	0.076	1.220	12.36
GraphSAGE	0.041	44.473	0.400	14.313	50.053	0.065	0.056	1.096	8.79
RelGT	0.035	43.992	0.326	14.267	48.922	0.065	0.054	1.030	6.71
RelGNN	0.037	44.461	<u>0.301</u>	14.230	48.767	0.065	0.054	1.027	6.79
RelAgent	0.033	<u>37.194</u>	0.386	13.949	**41.765**	0.064	<u>0.035</u>	**0.934**	**3.00**
RT-P+fine-tuning	0.0348	**32.7**	0.410	**13.9**	<u>43.0</u>	<u>0.0635</u>	0.0403	<u>0.953</u>	<u>3.57</u>
Foundation Model									
Closed-Source									
KumoRFM-1	0.035	58.231	0.417	16.161	55.254	0.065	0.040	1.102	9.50
KumoRFM-2 (orig.)*	0.034	43.293	0.433	<u>13.921</u>	46.992	0.064	**0.034**	0.986	5.14
KumoRFM-2 (repro.)*	0.033	41.974	0.433	14.627	45.352	0.065	0.043	1.015	6.50
TabPFN-Rel	**0.0311**	39.8	0.413	14.4	46.8	0.0679	0.0605	1.051	6.50
Open-source									
Griffin	0.050	78.232	0.463	35.590	53.214	0.092	0.151	1.737	16.71
RT	0.058	73.999	0.455	18.802	57.996	0.110	0.089	1.543	16.57
RT-J	0.0463	58.840	**0.0758**	17.360	60.630	**0.0584**	0.0581	0.957	10.29
RandGNN	0.038	53.370	0.434	17.580	66.130	0.071	0.065	1.248	14.14
TabPFN-Rel-OSS	<u>0.0314</u>	42.6	0.386	14.4	47.8	0.068	0.0614	1.058	7.29
RDBLearn	0.033	43.565	0.378	14.504	47.038	0.067	0.063	1.066	7.43

* See comment from Table 5.

F Figure 2 Setup Details and Extended Empirical Results

For our initial source of single tables, we select binary classification and regression tasks from BeyondArena (Purucker et al., 2026), Fraud Dataset Benchmark (FDB) (Grover et al., 2022), and OpenML (Vanschoren et al., 2014). The origins of specific tasks are listed in Table 7. To better align with realistic RDBs, we consider tasks that have temporal splits, e.g., the "Temporal" category within BeyondArena. Large datasets (*home_credit_default_stability*) and datasets without categorical columns (*maps_router_eta*) are excluded. To induce RDBs from these tables, we select at most 5 categorical columns as FKs following the procedure described in Section 4.

In Table 7 we present the full results that contributed to Figure 2, including performance from TabPFN-v3 (applied only on the original single-table data) (Grinsztajn et al., 2026), RDBLearn (Xu et al., 2026), KumoRFM-2 (Hudovernik et al., 2026), and RT-J (Ranjan et al., 2026b). Because BeyondArena has no validation set, and to standardize comparisons, we simply execute RDBLearn under the fixed setting of hop $H = 2$ along with just a single TFM selected as TabPFN-v3. For KumoRFM-2, we also use the same default hop value $H = 2$ and leave the inference-time configuration unchanged. RT-J uses the default inference configuration as well. Regarding hardware, we ran each method on a single GPU. TabPFN-v3, RDBLearn, and RT-J were run on a single Nvidia RTX 5090 (32G). Although the dataset scale is the same, KumoRFM-2 requires more GPU memory. To fulfill this requirement, we ran KumoRFM-2 on a single Nvidia L20 (48G). However, there were still four tasks raising OOM errors, noted in Table 7 and Figure 2.

Table 7: Full results corresponding to Figure 2, where we examin RDBs induced from single tables. Classification reports ROC-AUC (↑). Regression reports R^2 (↑). Best result per row in **bold**; second best underlined. Avg (excluding OOM) averages only over tasks where every method reports a value; Avg (all) averages over all tasks (KumoRFM-2 omitted due to OOM entries).

Task	Origin	TabPFN-v3	RDBLearn	KumoRFM-2	RT-J
Classification (ROC-AUC ↑)					
acquire_valued_shoppers_challenge	BeyondArena	**0.6245**	0.5967	OOM	0.5730
anes_voting_2026	BeyondArena	**0.8620**	0.8308	OOM	0.5445
bank_marketing_temporal_month	OpenML	**0.8856**	0.8769	0.8568	0.6384
electricity_temporal	OpenML	**0.9542**	0.9232	0.8752	0.7326
hotel_booking_demand	BeyondArena	**0.8059**	0.7712	0.6645	0.5895
kick	BeyondArena	**0.8220**	0.8001	0.7611	0.5290
kickstarter	BeyondArena	**0.9101**	0.9094	0.8354	0.7705
lending_club	BeyondArena	**0.6767**	0.6633	0.6311	0.5143
sparknov_temporal	FDB	**0.9809**	0.9398	0.9079	0.7500
Avg (excluding OOM)		***0.8622***	*0.8406*	*0.7903*	*0.6463*
Avg (all)		***0.8358***	*0.8124*	–	*0.6269*
Regression (R^2 ↑)					
california_house_prices_2020	BeyondArena	**0.9351**	0.9290	0.9091	0.5153
climate_model_weather_forecasting	BeyondArena	**0.8339**	**0.8339**	0.7879	0.6271
coffee_rating_prediction	BeyondArena	**0.3390**	0.2994	0.0928	0.0661
cooking_time	BeyondArena	**0.4549**	0.4477	0.3617	0.1948
delivery_eta	BeyondArena	**0.4375**	0.2757	0.4034	0.0192
garments_worker_productivity	BeyondArena	0.3686	**0.3734**	0.3011	−0.2398
mercedes_benz_greener_manufacturing	BeyondArena	**0.5519**	0.5453	OOM	0.0134
rossmann_store_sales	BeyondArena	**0.8161**	0.8146	0.5957	0.5895
sberbank_housing_market_forecasting	BeyondArena	**0.6283**	0.5829	OOM	0.1226
sf_permit_time	BeyondArena	0.0610	0.0479	**0.1669**	0.0419
Avg (excl. OOM)		***0.5308***	*0.5027*	*0.4523*	*0.2268*
Avg (all)		***0.5426***	*0.5150*	–	*0.1950*

G Figure 3 Setup Details and Extended Results

This section describes the detailed setup for experimenting with the interaction terms discussed in Sections 6 and 7, particularly as they relate to Figure 3.

G.1 Synthetic RDB and Interaction Experiment - Figure 3 (left)

Data generation. Each synthetic RDB contains a target table (4,000 rows: 3,200 for training and 800 for testing; two columns: PK and a label column), one child table, and nothing else. Each RDB links every target row to k rows in the child table $\boldsymbol{T}^2$ via PK-FK pairs, whose columns contain the two signal columns $\boldsymbol{a}, \boldsymbol{b}$ among a growing number of columns. We set $k \sim \text{Uniform}\{2, \ldots, 20\}$. Signal columns for the interaction term in the child table are synthesized with Bernoulli(0.3) for binary variants or $\mathcal{N}(0, 1)$ for Gaussian ones, alongside with $d - 2$ additional columns drawn i.i.d. from the same family, for total widths $d \in \{8, 16, 32, 64, 96\}$.

Label synthesis. We instantiate four labeling functions to cover a variety of linear and nonlinear relationships. For each row in the task table where PK$= i$, let $\mathcal{R}_i = \{j \in \boldsymbol{T}^2 \text{ where FK}_j = i\}$. Labels are generated by aggregating the row-level interaction with different aggregation functions, including `mean` ($y_i = \text{mean}_{j \in \mathcal{R}_i} a_j b_j$), `thresholded count`: ($y_i = \text{count}_{j \in \mathcal{R}_i}(a_j b_j > 1)$), `parity` ($y_i = \text{count}_{j \in \mathcal{R}_i}(a_j \oplus b_j)$), alongside the pairing-free control `prod_agg` ($y_i = \text{mean}_{\mathcal{R}_i}(\boldsymbol{a}) \cdot \text{mean}_{\mathcal{R}_i}(\boldsymbol{b})$). The latter is favorable to RDBLearn since g_{pf} already computes this form of interaction between aggregations, as mentioned in Section 6. Meanwhile, the other aggregation choices are not favorable to RDBLearn since they are not computed by any single aggregation (although combinations may approximate them). For different labeling rules, we experimented with different data column distributions as follows:

- `parity`, `mean`, and `prod_agg` with binary columns
- `count`$(ab > 1)$ with Gaussian columns.

Together these selections give us four different tasks as listed in Figure 3 (*left*).

Experiment models. We compare RDBLearn against the top non-TFM-based RDB FMs, KumoRFM-2 (Hudovernik et al., 2026) and RT-J (Ranjan et al., 2026b), all in-context with no task-specific training or fine-tuning. RDBLearn simply uses $H = 2$ to match the KumoRFM-2 default setting, run from a local deployment of the published model. RT-J uses the public `rt-j` checkpoint zero-shot. For both KumoRFM-2 and RT-J we first validated our implementation pipeline by reproducing performance numbers from the respective published papers.

G.2 Quasi-Real World RDB and Interaction Experiment - Figure 3 (right)

We adopt the rel-trial database (Robinson et al., 2024) and constituent tasks for the quasi-real world interaction experiment. Labels are constructed as temporally-valid mean values of the row-level product of two existing child-table columns. We do not modify the rel-trial schema; instead we pick two numeric/categorical columns $(\boldsymbol{a}, \boldsymbol{b})$ co-located in one child table of the task entity, and define each task label as the mean of the row-level product $a_i b_i$ over the linked child rows, computed strictly over rows dated before the task timestamp. The same per-row cutoff column is consumed by all three models, so no model can access post-cutoff rows during labeling or inference. Column pairs touching value columns whose measurement unit varies across rows (`param_value` and the confidence-interval limits, whose unit depends on a sibling code column) are excluded *a priori* and outcome-blind, leaving nine child-table pairs. These pairs form the x-axis labels of Figure 3(*right*).

We repeat the above experiment, this time explicitly adding an additional child-table column that formed from this interaction term. This is meant to reflect a hypothetical scenario whereby an informed agent has supplied this product term as an additional useful feature. We then explore the degree to which each RDB FM can effectively utilize this column to improve predictive performance. From Figure 3(*right*) we can see that RDBLearn is able to reliably exploit this added information, more so than the other models.

H Details of Real-World Agentic Experiment

Following the pointer at the end of Section 7, this experiment represents a preliminary proof-of-concept, not an exhaustive exploration of possibilities. And certainly there are more principled ways of incorporating agents within a TFM-based paradigm that what we have done. We also stress one

additional point upfront: Even as agents are often able to generate useful predictive features, this typically requires extensive domain knowledge and/or meta-data associated with a given RDB (and this introduces new pathways for leakage or memorization, as discussed in Appendix D, that may artificially inflate performance). In contrast, TFM-based models like RDBLearn and TabPFN-Rel can in principle operate out of the box in new domains without such knowledge, providing privacy and greater flexibility.

H.1 Experiment setup

TFM-based methods use agent-friendly SQL-based RDB harnesses, where SQL queries are automatically generated to aggregate features from RDBs. These queries are agent-readable and might provide the agent an alternative way to understand the predictive modeling task as well as the RDB structure. Based on this assumption, we prompt Cursor Grok 4.5 under the default Cursor 3.6.31 harness. We provide a single prompt together with the raw training set and the corresponding features generated by RDBLearn's encoder g_{pf}. Note that we do not allow the agent to iterate (e.g., by using a loop). Therefore, the feature generation is a single forward pass. This is substantially different from the SOTA RelAgent (Huang et al., 2026), which trains numerous fully-supervised models (on a per-task basis) across multiple iterations of a multi-agent system.

The prompts for all tasks are the same except for the task name. As a representative example, the rel-trial prompt is as follows:

Prompt for Cursor Grok 4.5 to generate RDB features

> If you look into what the rel-trial/study outcome tasks are, understand what the goal is, then look into DFS-depth 4 (including hop 1-4) SQLs. Can you analyze what features are useless and what new useful feature column you can add to the flattened table by SQL?"

H.2 Leakage Check

In our experiment, only the training set is exposed to the agent. To investigate whether there is data leakage, we manually look into the trace and prompt the agent to self-check. As a result, two features for rel-event tasks are found aggregating future (the same week) data. They are *sql_hist_ignore_rate* and *sql_hist_repeat_rate*, which calculate past-week label statistics. Although the target predictive entity is excluded from the calculation, when this predictive entity correlates with a mid-week date, same-week future label information can still be included. (For reference, performance results in Table 8 below exclude these two features.) Beyond this, inspecting every tool call in the agent trace, we did not observe any behavior that would indicate retrieval/access of validation or training set information. The full traces are provided in `https://anonymous.4open.science/status/iclr-submission-39479-35EC/`.

H.3 Results

Table 8 presents the performance comparison between the agent-feature-enhanced RDBLearn (based on TabPFN-v3.5 (Jäger et al., 2026)) and a SOTA multi-agent system, RelAgent (Huang et al., 2026). We note that even without domain knowledge and metadata information, RDBLearn already outperforms RelAgent (Qwen3-32B). And when enhanced with agent-generated features, its performance is significantly boosted to be on-par with RelAgent (GPT-5.2). Critically though, RelAgent is *fully-supervised*, training many per-task models, whereas RDBLearn remains a training-free FM enhanced with just a single batch of agent-generated features and no iterations.

I Technical Proofs

I.1 Proof of Proposition 5.2

Generalizing from Nagler (2023), the proof here is relatively straightforward once the RDB-dependent setup is specified through Definition 5.1, the assumption that (2) holds, and that $p\big(y_i|\mathcal{G}_H(\boldsymbol{x}_i)\big) = p\big(y_i|\boldsymbol{z}_i\big)\ \forall i = 0, \dots, n$. The first and third assumption alone imply that a given

Table 8: RDBLearn+Agent results on Relbench v1 binary classification tasks (AUROC↑). Bold = best per column, underline = second best.

Dataset	Task	RelAgent (GPT-5.2)	RelAgent (Qwen3-32B)	RDBLearn	RDBLearn+Agent
rel-amazon	item-churn	**0.8284**	0.7992	0.8270	<u>0.8273</u>
	user-churn	**0.7078**	0.6922	<u>0.6973</u>	0.6968
rel-avito	user-clicks	0.6836	0.6510	**0.7050**	<u>0.6994</u>
	user-visits	**0.6779**	0.6462	<u>0.6714</u>	0.6661
rel-event	user-ignore	**0.8725**	0.7591	0.7446	<u>0.7873</u>
	user-repeat	**0.7820**	0.7473	0.7646	<u>0.7753</u>
rel-f1	driver-dnf	<u>0.7834</u>	0.5767	0.7263	**0.8351**
	driver-top3	<u>0.8523</u>	0.5804	0.7975	**0.9204**
rel-hm	user-churn	**0.7107**	<u>0.7062</u>	0.6895	0.6953
rel-stack	user-badge	**0.8842**	0.6530	<u>0.8678</u>	0.8514
	user-engagement	<u>0.9041</u>	0.8911	**0.9051**	0.8982
rel-trial	study-outcome	0.7186	0.7136	<u>0.7265</u>	**0.7569**
AVG (12 tasks)		<u>0.7838</u>	0.7013	0.7602	**0.7841**

label y_i is independent of others when conditioned on the corresponding encoder outputs, i.e., $p(y_i, y_j|\boldsymbol{z}_i, \boldsymbol{z}_j) = p(y_i|\boldsymbol{z}_i)p(y_j|\boldsymbol{z}_j)\ \forall i, j$. Hence we may also infer that for any labeled sample, the conditional $p(y|\boldsymbol{z})$, as opposed to merely $p\big(y|\mathcal{G}_H(\boldsymbol{x})\big)$, is stable across n, even while encoder representations $\boldsymbol{z}$ need not be.

Now consider any sequence pairs $p, \check{p} \in \mathcal{P}$, along with their associated label generators $p(y|\boldsymbol{z})$ and $\check{p}(y|\boldsymbol{z})$ respectively, satisfying $p(y|\boldsymbol{z}) = \check{p}(y|\boldsymbol{z})$ whenever $\|\boldsymbol{z}_{\text{test}} - \boldsymbol{z}\| < \epsilon$ for any fixed $\epsilon > 0$. For all such pairs, we have from (2) (bias part) that

$$\begin{aligned} p(y_{\text{test}}|\boldsymbol{z}_{\text{test}}) - \mathbb{E}_{\mathcal{D}_n\sim p_n}\left[q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right] \\ - \left(\check{p}(y_{\text{test}}|\boldsymbol{z}_{\text{test}}) - \mathbb{E}_{\mathcal{D}_n\sim \check{p}_n}\left[q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right]\right) \\ \mathbb{E}_{\check{\mathcal{D}}_n\sim \check{p}_n}\left[q_\theta\big(y_{\text{test}} \mid \{\check{\boldsymbol{z}}_i, \check{y}_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right] - \mathbb{E}_{\mathcal{D}_n\sim p_n}\left[q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right] \ \to\ 0 \end{aligned} \tag{13}$$

as $n \to \infty$, where $\{\boldsymbol{z}_{\text{test}}, y_{\text{test}}\}$ is treated as a fixed evaluation point. Combining with (2) (variance part) we then may conclude that

$$\left|q_\theta\big(y_{\text{test}} \mid \{\check{\boldsymbol{z}}_i, \check{y}_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big) - q_\theta\big(y_{\text{test}} \mid \{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big)\right| \ \to\ 0. \tag{14}$$

From standard diagonalization arguments, the same holds true for some $\epsilon_n \to 0$ since no assumptions were placed on the original ϵ, therefore (3) is established. ■

I.2 Proof of Proposition A.3

We show this result through an interpretable counter-example. Note that there is nothing particularly special about the configuration we introduce for this purpose; rather it merely represents a simplified case sufficient for illustration purposes and completing the proof. As we will see, the core ingredient is the existence of at least some degree of nonlinear interaction between relational effects within $\mathcal{G}_H(\boldsymbol{x})$ and the original tabular features $\boldsymbol{x}$.

Basic Setup. For any $\varepsilon \in (0, 1)$ and $C > 0$ we define mappings $f_1' : \mathbb{R}^d \to \mathbb{R}^2$, $f_2' : \mathbb{G}_H \to \mathbb{R}^2$, and $\Psi' : \mathbb{R}^2 \to \mathbb{R}$ as

$$f_1' : \begin{bmatrix} u_1 \\ 0 \end{bmatrix} = \begin{bmatrix} \boldsymbol{\omega}^\top \widetilde{\boldsymbol{x}}^{\text{real}} \\ 0 \end{bmatrix} + \begin{bmatrix} \beta \\ 0 \end{bmatrix}, \qquad f_2' : \begin{bmatrix} 0 \\ u_2 \end{bmatrix} = \begin{bmatrix} 0 \\ \pi\left(\mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}\right) \end{bmatrix}, \tag{15}$$

$$\Psi' :\ y = wv_1\text{sign}[v_2] + v_1 + v_2,$$

where $\boldsymbol{\omega} \in \mathbb{R}^{d'}$, $\beta \in \mathbb{R}$, and $w \in \mathbb{R}$ are constants in general position, and $\pi : \mathbb{G}_H \to \mathbb{R}$. For the latter, we choose π such that under p_τ we have

$$\text{Prob}\left[\pi\big(\mathcal{G}_H(\boldsymbol{x}_{\text{test}})\backslash\boldsymbol{x}_{\text{test}}\big) > 0\right] \ =\ \text{Prob}\left[\pi\big(\mathcal{G}_H(\boldsymbol{x}_{\text{test}})\backslash\boldsymbol{x}_{\text{test}}\big) < 0\right] \ =\ 0.5. \tag{16}$$

Additionally, underpinning (15) we use $\mathbb{G}_H$ to denote the relational neighborhood space of $\mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}$, while $\{v_1, v_2\}$ and $\{u_1, u_2\}$ are the variables adopted to express the relevant input and output spaces, respectively. Combining these specifications with (9) we arrive at the label generation function

$$y = w(1-\tau)\tau u_1\text{sign}[u_2] + (1-\tau)u_1 + \tau u_2, \tag{17}$$

where we observe nonlinear coupling between features derived from $\boldsymbol{x}$ (expressed through u_1) and $\mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}$ (expressed through u_2).

Handling Equalities in (11). To establish the first equality of (11), we observe that by construction $y = \boldsymbol{\omega}^\top \widetilde{\boldsymbol{x}}^{\text{real}} + \beta$ when $\tau = 0$. However, we may infer that

$$\min_{f_\theta^{\text{tab}} \in \mathcal{F}} \mathbb{E}_{\mathcal{D}_0 \sim p_0} \Big[\, \big| y_{\text{test}} - \hat{y}_{\text{test}}^{\text{tab}} \big| \, \Big] \;=\; \min_{f_\theta^{\text{tab}} \in \mathcal{F}} \mathbb{E}_{\mathcal{D}_0 \sim p_0} \Big[\, \big| \boldsymbol{\omega}^\top \widetilde{\boldsymbol{x}}_{\text{test}}^{\text{real}} + \beta - f_\theta^{\text{tab}}\big(\{\boldsymbol{x}_i, y_i\}_{i=1}^n, \boldsymbol{x}_{\text{test}}\big) \big| \, \Big] \;=\; 0, \tag{18}$$

where p_0 follows Definition A.1 with $\tau = 0$. Given the assumption that $n > d'$, the equality with zero in (18) occurs by straightforward extension of Xu et al. (2026)[Proposition 4.2], which is not operationally effected by padding $\boldsymbol{x}_{\text{test}}$ and each $\boldsymbol{x}_i$ with additional categorical elements. Meanwhile, the second equality in (11) stems from the fact that $y = \pi\big(\mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}\big)$ and our encoder constraint set $\bar{g}$ allows for modeling an arbitrary function π over the relational neighborhood excluding $\boldsymbol{x}$. Hence when $\tau = 1$ we trivially have that

$$\begin{aligned} &\min_{f_\theta^{\text{rdb}} \in \mathcal{F},\; g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_1 \sim p_1} \Big[\, \big| y_{\text{test}} - \hat{y}_{\text{test}}^{\text{rdb}} \big| \, \Big] \\ &\qquad\equiv \min_{f_\theta^{\text{rdb}} \in \mathcal{F},\; g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_1 \sim p_1} \Big[\, \big| \pi\big(\mathcal{G}_H(\boldsymbol{x}_{\text{test}})\backslash\boldsymbol{x}_{\text{test}}\big) - f_\theta^{\text{rdb}}\big(\{\boldsymbol{z}_i, y_i\}_{i=1}^n, \boldsymbol{z}_{\text{test}}\big) \big| \, \Big] \;=\; 0. \end{aligned} \tag{19}$$

Handling the Inequality in (11). Lastly, we handle the third line of (11), namely, the inequality that results for τ values bounded away from zero and one. We establish this bound via several steps as follows. We also remark upfront that α can be absorbed into the specifications of $\hat{y}_{\text{test}}^{\text{tab}}$ and $\hat{y}_{\text{test}}^{\text{tab}}$ without loss of generality.

First, because of the constraint $g_\phi \in \bar{g}$, the encoder cannot reliable reconstruct $\boldsymbol{x}_{\text{test}}^{\text{real}}$; at most we can guarantee reconstruction of $d - 1$ coordinates per Definition A.2. Moreover, it is not possible to reconstruct the missing coordinate from the others with a continuous mapping from $\mathcal{F}$, as doing so would violate the assumption from Definition A.1 that any $p(\boldsymbol{x}^{\text{real}})$ is bounded from above (i.e., if we could reconstruct this coordinate with a continuous mapping, it would imply that the data lies on a low-dimensional manifold within $\mathbb{R}^d$, which by necessity would have infinite density).

Given these considerations, we can again invoke Xu et al. (2026)[Proposition 4.2] to infer that there will be non-negligible error (in expectation over the generation process) for any attempt to predict $u_{1,\text{test}}$ from available context $\{\boldsymbol{z}_i, y_i\}_{i=1}^n$ and $\boldsymbol{z}_{\text{test}}$. Intuitively, this is because the support pattern for extracting $\widetilde{\boldsymbol{x}}_{\text{test}}^{\text{real}}$ from $\boldsymbol{x}_{\text{test}}^{\text{real}}$ varies uniformly across random draws from p_τ, and so in expectation there will always be cases where needed information is missing. From this we may conclude that

$$\min_{f_\theta^{\text{rdb}} \in \mathcal{F},\; g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau} \big[\, \big| u_{1,\text{test}} - \hat{y}_{\text{test}}^{\text{rdb}} \big| \big] \;=\; \delta_\tau \tag{20}$$

for any $\tau \in [\varepsilon, 1 - \varepsilon]$ and some $\delta_\tau > 0$. But this will also remain true when conditioning on the corresponding $u_{2,\text{test}}$. This is because $u_{2,\text{test}}$ is merely a function of other rows within $\boldsymbol{X}$, and conditioning on any or all rows (or functions thereof), excluding the current $\boldsymbol{x}_{\text{test}}$, will not remove the prediction error, since these rows are generated IID by Definition A.1. And so we have that

$$\min_{f_\theta^{\text{rdb}} \in \mathcal{F},\; g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau} \big[\, \big| u_{1,\text{test}} - \hat{y}_{\text{test}}^{\text{rdb}} \big| \;\big|\; u_{2,\text{test}} \geq \lambda \, \big] \;=\; \delta'_\tau \tag{21}$$

for some possibly new $\delta'_\tau > 0$; an analogous error term is bounded away from zero when conditioning on the event $u_{2,\text{test}} < 0$. And therefore from the tower rule of conditional expectations, we also have

$$\min_{f_\theta^{\text{rdb}},\; g_\phi \in \bar{g}} \mathbb{E}_{\mathcal{D}_\tau \sim p_\tau} \big[\, \big| u_{1,\text{test}} \text{sign}\big[u_{2,\text{test}}\big] - \hat{y}_{\text{test}}^{\text{rdb}} \big| \, \big] \;=\; \delta''_\tau \;>\; 0, \tag{22}$$

which becomes $w(1-\tau)\tau\delta''_\tau$ once we reintroduce scale factors. If we choose $w = \frac{C}{(1-\varepsilon)\varepsilon\delta^*}$ with $\delta^* := \inf_{\tau \in [\varepsilon, (1-\varepsilon)]} \delta''_\tau$, then

$$w(1-\tau)\tau\delta''_\tau \;\;\geq\;\; \frac{C(1-\tau)\tau}{(1-\varepsilon)\varepsilon} \;\;\geq\;\; C \quad \forall \tau \in [\varepsilon, (1-\varepsilon)]. \tag{23}$$

Finally, the remaining factors $(1-\tau)u_{1,\text{test}}$, $\tau u_{2,\text{test}}$, and $\hat{y}_{\text{test}}^{\text{tab}}$ from within the last expectation of (11) cannot appreciably improve the lower bound of C. Note that the factor $\tau u_{2,\text{test}}$ can be absorbed into

$\hat{y}^{\text{rdb}}_{\text{test}}$ without loss of generality; they share the same input range while the former is a pre-specified function and the latter is arbitrary.

Analogously, w.l.o.g. $(1-\tau)u_{1,\text{test}}$ can be absorbed into $\hat{y}^{\text{tab}}_{\text{test}}$. But the flexibility from $\hat{y}^{\text{tab}}_{\text{test}}$ itself is inconsequential to altering the bound. This is because $u_{1,\text{test}} \perp\!\!\!\perp u_{2,\text{test}}$, given that $\boldsymbol{x}^{\text{real}}_{\text{test}} \perp\!\!\!\perp \mathcal{G}_H(\boldsymbol{x})\backslash\boldsymbol{x}$ per Definition A.1, and so reducing the error approximating the interaction term $u_{1,\text{test}}\text{sign}[u_{2,\text{test}}]$ is not possible without access to $u_{2,\text{test}}$ (noting that $u_{2,\text{test}}$ has a median of zero). ■

I.3 Proof of Proposition B.1

High-Level Intuition. To illustrate the gist of this result, assume that the ground-truth label y_{test} for any $\boldsymbol{x}_{\text{test}}$ is given by a parity check on available labels in $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$. This labeling rule is impossible to determine from $y_H(\boldsymbol{x}_{\text{test}})$ alone, since for any node $\boldsymbol{x}'$ associated with a label $y' \in y_H(\boldsymbol{x}_{\text{test}})$, the corresponding set $y_H(\boldsymbol{x}')$ need not be fully observed, i.e., the overlap between $y_H(\boldsymbol{x}')$ (with a known label for $\boldsymbol{x}'$) and $y_H(\boldsymbol{x}_{\text{test}})$ (with an unknown label for $\boldsymbol{x}_{\text{test}}$) is only partial. Hence the parity rule is not identifiable, as even a single missing label can flip the result. This disrupts inductive inference from the information available within $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$, which is all a fixed FM encoder has access too. In contrast, if per-dataset training is allowed over RDB instances formed via a given parity-based labeling rule, it is straightforward to learn the associated mapping on a dataset-by-dataset or task-by-task basis. We remark that, although the proof below relies on parity checks for illustrative convenience, the underlying phenomena is emblematic of quite broad scenarios where diverse heterophilic (as opposed to homophilic) relationships are dominant.

Basic Setup. The basic proof follows from the construction of a suitable joint distribution over context $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$ and test label y_{test}. Meanwhile, by assumption we have that $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ is fixed and given. To begin, we introduce a latent indicator variable $z \in \{0,1\}$ that selects among two deterministic generative processes for $\{\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}\}$. In this way, all randomness in the joint distribution $p(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}} \mid \mathcal{G}^*_H(\boldsymbol{x}_{\text{test}}))$ will ultimately reduce to randomness from $p(z)$. We emphasize that this is not meant to exclude broader sources of randomness that naturally describe RDB content; instead, we are merely distilling down to a minimal design case that nonetheless still allows us to establish the proposition. Other latent sources of variation associated with real-world RDBs then become largely orthogonal to our result and need not be explicitly specified.

Label Assumptions. We assume that all labels are binary, namely, $y \in \{0,1\}$. While generalization to broader label domains $\mathcal{Y}$ is feasible, it introduces superfluous complexity given that the binary case is already sufficient for establishing the proof. For any $y_H(\boldsymbol{x}) \neq \emptyset$, where $y_H(\boldsymbol{x})$ indicates the set of all past labels present within $\mathcal{G}^*_H(\boldsymbol{x})$ for some arbitrary $\boldsymbol{x}$, we define conditional label generation via a parity function and its negative given by

$$\begin{aligned} \pi\big[y_H(\boldsymbol{x})\ ;\ z=1\big] &:= \left[\sum_{y\in y_H(\boldsymbol{x})} y\right] \bmod 2 \\ \pi\big[y_H(\boldsymbol{x})\ ;\ z=0\big] &:= \neg\pi\big[y_H(\boldsymbol{x}); z=1\big], \end{aligned} \tag{24}$$

both of which output a label $y \in \{0,1\}$.

Context Assumptions. We specify a function $h\big(\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}}); z\big)$ for context generation. To ensure no conflicting labels, for $z=1$ we define this h such that $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$ has label assignments consistent with both $y_H(\boldsymbol{x}_{\text{test}})$ and the label assignment rule $\pi\big[\,\cdot\ ;\ z=1\big]$ from (24). This will always be achievable by introducing new appropriately-labeled neighbors to specific nodes of $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ associated with elements of $y_H(\boldsymbol{x}_{\text{test}})$.

To establish this possibility, assume an arbitrary trial arrangement of node features, edges, and labels within $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$, but following the typing of $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ for consistency. Then consider any $y' \in y_H(\boldsymbol{x}_{\text{test}})$ with corresponding $\boldsymbol{x}'$ extracted from the associated row of $\boldsymbol{X} \equiv \boldsymbol{T}^K$ (per original definitions in the main text). For consistency with the labeling rule (24) with $z=1$, we require that $y' = \pi\big[y_H(\boldsymbol{x}')\ ;\ z=1\big]$, which may not initially be the case since labeling assignments were thus far

assumed arbitrary. However, for any node where this is not the case, we can add a single additional labeled neighbor to $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$, with an earlier time-stamp and a label selected to flip the parity check, such that we obtain the necessary equality.

Such a neighbor can be added in multiple ways given that $H \geq 2$ (note that if $H = 1$ then $|y_H(\boldsymbol{x}_{\text{test}})| = 0$, i.e., no labels within one hop, which is not allowed by assumption). First, assume $H = 2$, if some table T^k exists with at least two FK columns pointing to the PK column associated with $T^K \equiv \boldsymbol{X}$, we can trivially add such a neighbor by adding a new labeled row to T^K and a new row to T^k that links the former to $\boldsymbol{x}' \in T^K$. Meanwhile, if such a T^k does not exist, we can simply add such a table to $\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})$. Moreover, when supplied with an earlier time-stamp, the added labeled node/row need not introduce any additional constraints vis-à-vis (24). With regard to the latter, if $\boldsymbol{x}''$ indicates the added node, we can directly enforce $y_H(\boldsymbol{x}'') = \emptyset$ as no nodes with earlier time-steps exist, and so y'' can be selected free of any constraint from (24), which would only apply if $y_H(\boldsymbol{x}'') \neq \emptyset$. Lastly, because this added label will by design be two hops from $\boldsymbol{x}'$ and three or more hops from any other labeled rows, it will not introduce new conflicts per (24) if $H = 2$. For $H > 2$ an analogous procedure can be applied pushing the new labeled node to H hops away from $\boldsymbol{x}'$ so as to achieve the same effect.

In an analogous manner we can also define h when $z = 0$, handling label consistency via the introduction of neighboring nodes to adjust parity accordingly. Again, we are granted this flexibility by assumption, since only $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})$ is fixed and given.

Finalizing Part 1. Per the assumptions stated above and the appropriate counting measure, it directly follows that $\int p\left(y_{\text{test}}|\mathcal{G}^*_{\text{test}}\right) \log p\left(y_{\text{test}}|\mathcal{G}^*_{\text{test}}\right) dy_{\text{test}} = 0$ trivially establishing that y_{test} is a deterministic function of $\mathcal{G}^*_{\text{test}}$. Furthermore, given $\mathcal{G}^*_{\text{test}} = \left\{\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), \mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right\}$ by definition, we also have that

$$
\begin{aligned}
&\mathbb{E}_{p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right)}\left[\mathbb{KL}\Big[p(y_{\text{test}}|\mathcal{G}^*_{\text{test}}) \;||\; q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big)\Big]\right] \\
&= \int p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right) \int p\left(y_{\text{test}}|\mathcal{G}^*_{\text{test}}\right) \log \frac{p\left(y_{\text{test}}|\mathcal{G}^*_{\text{test}}\right)}{q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big)} dy_{\text{test}} d\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}) \\
&= -\int p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right) \int p\left(y_{\text{test}}|\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), \mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right) \log q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big) dy_{\text{test}} d\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}) \\
&= -\int\int p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right) \log q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big) dy_{\text{test}} d\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}) \qquad (25)
\end{aligned}
$$

for any fixed g_ϕ and q_θ. By construction, $p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right) \neq 0$ at only two points associated with $z = 0$ and $z = 1$. Let $y^{(0)}_{\text{test}}$ and $y^{(1)}_{\text{test}}$ denote the value of y_{test} at these two points. Moreover, if we assume that z is drawn uniformly over $\{0, 1\}$, we have $\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}}) \perp\!\!\!\perp z$ and so we may conclude from (25) that

$$
\begin{aligned}
&\mathbb{E}_{p\left(\widetilde{\mathcal{G}}^*_H(\boldsymbol{x}_{\text{test}})|\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\right)}\left[\mathbb{KL}\Big[p(y_{\text{test}}|\mathcal{G}_{\text{test}}) \;||\; q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big)\Big]\right] \\
&\qquad = \; -p(z=0) \log q_\theta\big(y^{(0)}_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big) \;-\; p(z=1) \log q_\theta\big(y^{(1)}_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big) \\
&\qquad \geq \; \mathbb{H}[z] \;=\; \log 2. \qquad (26)
\end{aligned}
$$

This lower bound is achievable iff $q_\theta\big(y_{\text{test}} \mid g_\phi\big[\mathcal{G}^*_H(\boldsymbol{x}_{\text{test}})\big]\big)$ is a uniform distribution over $\{0, 1\}$.

Finalizing Part 2. Assume that some parameter-free encoder g computes the element-wise sum of $y_H(\boldsymbol{x})$; other arbitrary features may be computed as well, but they are are ultimately treated as spurious features in what follows. Given this setup, the prediction head q_θ need only learn to differentiate even and odd sums associated with $y_H(\boldsymbol{x})$, while ignoring all other feature columns. For example, if $z = 1$ governs the generative process, then odd values of any sum over $y_H(\boldsymbol{x})$ will be mapped to 1, even values to 0. Such functions can be directly represented by ReLU networks, and while not a requirement for this proof, are even efficiently learnable with additional assumptions (Daniely & Malach, 2020). ■

I.4 Proof of Proposition B.2

High-Level Intuition. The proof follows through modifications applied to the original parity setup from Appendix I.3. For illustrative simplicity, assume a single informative binary RDB feature column, whereby ground-truth labels are a deterministic function, akin to a parity check, on the values of all features from this column contained within $\mathcal{G}_H^*(\boldsymbol{x}_{\text{test}})$. All other RDB columns are also binary but uninformative. Inclusion of $y_H(\boldsymbol{x}_{\text{test}})$ within $\mathcal{G}_H^*(\boldsymbol{x}_{\text{test}})$ does not allow for differentiating which column is informative, for analogous reasons to those presented in establishing Proposition B.1. In brief, for any node $\boldsymbol{x}'$ associated with a label $y' \in y_H(\boldsymbol{x}_{\text{test}})$, the corresponding set $y_H(\boldsymbol{x}')$ *and its associated features* need not be fully observed, and so a parity function on *any* column can produce an arbitrary distribution of $y_H(\boldsymbol{x}_{\text{test}})$. Consequently, the parity rule is not identifiable regardless of whether $y_H(\boldsymbol{x}_{\text{test}})$ is present or not as an encoder input. Note that this line of reasoning is not restricted to binary feature columns, as we can quantize continuous or categorical columns into binary ones in forming deterministic labeling functions for analysis purposes. We now turn to the actual proof.

Basic Setup. As in the proof of Proposition B.1, all randomness is encapsulated by a binary latent variable z that selects between two deterministic generative processes, now over $\{\widetilde{\mathcal{G}}_H^*(\boldsymbol{x}_{\text{test}}), y_H(\boldsymbol{x}_{\text{test}}), y_{\text{test}}\}$, i.e., only $\mathcal{G}_H(\boldsymbol{x}_{\text{test}})$ as opposed to $\mathcal{G}_H^*(\boldsymbol{x}_{\text{test}})$ is assumed fixed/given. When granted $\mathcal{G}_H(\boldsymbol{x})$ for some arbitrary $\boldsymbol{x}$, we specify the conditional label generation functions

$$\begin{aligned}
\pi\big[\mathcal{G}_H(\boldsymbol{x})\,;\; z=1\big] \quad &:= \quad \left[\sum_{\{i,j,k\}\in S_H^{(1)}(\boldsymbol{x})} \beta\left(t_{ij}^k\right)\right] \bmod 2 \\
\pi\big[\mathcal{G}_H(\boldsymbol{x})\,;\; z=0\big] \quad &:= \quad \left[\sum_{\{i,j,k\}\in S_H^{(0)}(\boldsymbol{x})} \beta\left(t_{ij}^k\right)\right] \bmod 2,
\end{aligned} \tag{27}$$

where $S_H^{(1)}(\boldsymbol{x})$ and $S_H^{(0)}(\boldsymbol{x})$, associated with $z=1$ and $z=0$ respectively, represent distinct support sets over RDB column elements reachable within H hops from $\boldsymbol{x}$ following PK-FK pairs. In this way any $\{i,j,k\} \in S_H^{(1)}(\boldsymbol{x})$ (or analogously, within $S_H^{(0)}(\boldsymbol{x})$) indicates the i-th element within a specific RDB column $\boldsymbol{t}_{:j}^k$, namely column j in table k. Meanwhile β denotes a mapping from each table element to a binary value (this discretization is by no means necessary, but it facilitates a simplified exposition that is sufficient for establishing the core result). These values are then subsequently combined via the stated parity functions in (27) to produce binary labels $y \in \{0,1\}$.

Establishing Part 1. By adopting (27) as the label generating function, we trivially achieve the conditions of part 1 in Proposition B.2. Specifically, with all randomness deferred to z, which explicitly fixes the support set, all labels default to a deterministic function of the chosen support set, either $S_H^{(0)}(\boldsymbol{x})$ or $S_H^{(1)}(\boldsymbol{x})$ for any $\boldsymbol{x}$.

Establishing Part 2. If we stipulate that z is drawn uniformly over $\{0,1\}$ independently of $\mathcal{G}_H(\boldsymbol{x}_{\text{test}})$, then we have

$$p\left(S = S_H^{(0)}(\boldsymbol{x}_{\text{test}}) \mid \mathcal{G}_H(\boldsymbol{x}_{\text{test}})\right) \;=\; p\left(S = S_H^{(1)}(\boldsymbol{x}_{\text{test}}) \mid \mathcal{G}_H(\boldsymbol{x}_{\text{test}})\right) \;=\; \tfrac{1}{2}. \tag{28}$$

However, by design we also have the flexibility to select the distribution over $\widetilde{\mathcal{G}}_H^*(\boldsymbol{x}_{\text{test}})$ such that $y_H(\boldsymbol{x}_{\text{test}})$ is invariant to z, even while y_{test} still depends on z through (27) applied at $\boldsymbol{x}_{\text{test}}$. To see this, note that each $y' \in y_H(\boldsymbol{x}_{\text{test}})$ is computed as $\pi[\mathcal{G}_H(\boldsymbol{x}'); z]$, where $\boldsymbol{x}'$ is the target table feature row associated with y'. Hence we only need enforce

$$\pi\big[\mathcal{G}_H(\boldsymbol{x}')\,;\; z=1\big] = \pi\big[\mathcal{G}_H(\boldsymbol{x}')\,;\; z=0\big] \quad \text{for each } \; y' \in y_H(\boldsymbol{x}_{\text{test}}). \tag{29}$$

But this will always be possible per the following: Suppose the equality from (29) does not initially hold. We may add an additional element to either $S_H^{(1)}(\boldsymbol{x}')$ or $S_H^{(0)}(\boldsymbol{x}')$ (but not both) such that the corresponding tuple produces $\beta(t_{ij}^k) = 1$ to flip the parity and induce equality. This is achievable without disrupting equality elsewhere by placing the added element within $\widetilde{\mathcal{G}}_H^*(\boldsymbol{x}_{\text{test}})$, at a location H

hops from $\boldsymbol{x}'$ and greater than H hops from other elements of $y_H(\boldsymbol{x}_{\text{test}})$. This placement mirrors an analogous placement strategy applied in Appendix I.3, only here labeled nodes are not required.

Lastly, if (29) holds, then

$$p\big(S|\mathcal{G}_H^*(\boldsymbol{x}_{\text{test}})\big) = p\big(S|\mathcal{G}_H(\boldsymbol{x}_{\text{test}}), y_H(\boldsymbol{x}_{\text{test}})\big) = p\big(S|\mathcal{G}_H(\boldsymbol{x}_{\text{test}})\big), \tag{30}$$

completing the demonstration of part 2. ■