# **REDELEX:** A Framework for Relational Deep Learning Exploration

Jakub Peleška (⊠) and Gustav Šír

Czech Technical University in Prague, Karlovo náměstí 13, Prague, 121 35, Czechia jakub.peleska@fel.cvut.cz,gustav.sir@cvut.cz

Abstract. Relational databases (RDBs) are widely regarded as the gold standard for storing structured information. Consequently, predictive tasks leveraging this data format hold significant application promise. Recently, Relational Deep Learning (RDL) has emerged as a novel paradigm wherein RDBs are conceptualized as graph structures, enabling the application of various graph neural architectures to effectively address these tasks. However, given its novelty, there is a lack of analysis into the relationships between the performance of various RDL models and the characteristics of the underlying RDBs. In this study, we present REDELEX—a comprehensive exploration frame-

work for evaluating RDL models of varying complexity on the most diverse collection of over 70 RDBs, which we make available to the community. Benchmarked alongside key representatives of classic methods, we confirm the generally superior performance of RDL while providing insights into the main factors shaping performance, including model complexity, database sizes and their structural properties.

Keywords: Relational Deep Learning  $\cdot$  Relational Databases  $\cdot$  Graph Neural Networks

### 1 Introduction

From their establishment [9], Relational Databases (RDBs) played a pivotal role in transforming our society into the current information age. Data stored as interconnected tables, safeguarded by integrity constraints, have proven to be an effective method for managing domain information. Consequently, RDBs still prevail today as a backbone of critical systems in a number of important domains ranging from healthcare [41] to government [28].

Although ubiquitous in modern application stacks, the data format of RDBs is deeply incompatible with classic Machine Learning (ML) workflows, which assume data in the standard form of fixed-size i.i.d. feature vectors, forming the common "tabular" learning format. Nevertheless, this assumption is clearly violated with the relationships between the differently-sized RDB tables. To address the discrepancy, the historically prevailing approach has been to turn the relational into the tabular format by means of "propositionalization" [25],

which is essentially a feature extraction routine where relational substructures get aggregated from the relations into the attributes (features) of the tabular format, upon which classical ML methods may then operate. Nevertheless, this comes at the cost of information loss during this preprocessing step.

Recently, building on advances in graph representation learning [19], deep learning models directly exploiting the relational structure of RDBs have started to gain traction [13,44,43,32], establishing the field of *Relational Deep Learning* (RDL) [15]. Following the "message-passing" principles of *Graph Neural Networks* (GNN; [42]), RDL models treat the structure of an RDB as a heterogeneous (temporal) graph, where individual table rows correspond to nodes, and edges are formed through integrity constraints set by the primary and foreign keys. Utilizing the graph representation then allows for the application of various GNNs, and their various extensions, with adapted message-passing schemes.

The generality and spread of RDBs allow for a broad spectrum of domain information to be stored, upon which a variety of predictive tasks can be formulated, each with unique aspects and qualities. This presents a challenge for establishing a broad enough benchmark to appropriately assess the general performance of RDL. Currently, the most prominent effort in this area is the recently proposed RELBENCH [33], which introduced the evaluation of RDL, albeit with a very limited scope of simple models and just five accessible datasets. However, the overarching domain of *relational learning* [14], currently ignored by the RDL community, has a rich history of working with the relational data format [30,12], including benchmarking of the propositionalization techniques [25]. Notably, this includes the CTU Relational Learning Repository [29] that historically collected more than 70 diverse RDBs.

Our aim in this paper is to provide a bridge between the communities of traditional (logic-based) relational learning [14] and the contemporary RDL [33] towards a more comprehensive evaluation of the diverse existing methods. To that aim, we introduce REDELEX—an experimental framework for developing and benchmarking diverse RDL architectures against classic methods over the most comprehensive collection of tasks and datasets to date. The implementation of the framework is readily available on GitHub.<sup>1</sup>

# 2 Background

In this paper, we experimentally explore learning from RDBs (Sec. 2.1) with GNN-based models (Sec. 2.2) resulting in the RDL methodology (Sec. 2.3).

#### 2.1 Relational Databases

Principles of RDBs are formally based on the *relational model* [10], which is grounded in relational logic [16]. This abstraction enables the definition of any database, regardless of specific software implementation, as a collection of n-ary relations, which are defined over the domains of their respective attributes,

<sup>&</sup>lt;sup>1</sup> https://github.com/jakubpeleska/ReDeLEx

managed by the Relational Database Management System (RDBMS) to ensure data consistency with the integrity constraints of the database schema. The key concepts to be used in this paper are as follows.

**Relational Database** A Relational Database (RDB)  $\mathcal{R}$  is defined as a finite set of relations  $R_1, R_2, \ldots, R_n$ . An instance of an RDB  $\mathcal{R}$  is implemented through a RDBMS, enabling to perform Structured Query Language (SQL; [5]) operations, rooted in relational algebra.

**Relation (Table)** Formally, an *n*-ary relation  $R_{/n}$  is a subset of the Cartesian product defined over the domains  $D_i$  of its *n* attributes  $A_i$  as  $R_{/n} \subseteq D_1 \times D_2 \times \cdots \times D_n$ , where  $D_i = \operatorname{dom}(A_i)$ . Each relation *R* consists of a heading (signature)  $R_{/n}$ , formed by the set of its attributes, and a body, formed by the values of the respective attributes, commonly represented as a *table*  $T_R$  of the relation *R*.

Attribute (Column) Attributes  $\mathcal{A}_R = \{A_1, \ldots, A_n\}$  define the terms of a relation  $R_{/n}$ , corresponding to the *columns* of the respective table  $T_R$ . Each attribute is a pair of the attribute's name and a *type*, constraining the domain of each attribute as  $\mathsf{dom}(A_i) \subseteq \mathsf{type}(D_i)$ . An attribute *value*  $a_i$  is then a specific valid value from the respective domain of the attribute  $A_i$ .

**Tuple (Row)** An n-tuple in a relation  $R_{/n}$  is a tuple of attribute values  $t_i = (a_1, a_2, \ldots, a_n)$ , where  $a_j$  represents the value of the attribute  $A_j$  in R. The relation can thus be defined extensionally by the *unordered* set of its tuples:  $R = \{t_1, t_2, \ldots, t_m\}$ , corresponding to the rows of the table  $T_R$ .

**Integrity constraints** In addition to the domain constraints  $\operatorname{dom}(A_i)$ , the most important integrity constraints are the primary and foreign keys. A primary key PK of a relation R is a minimal subset of its attributes  $R[PK] \subseteq \mathcal{A}_{\mathcal{R}}$  that uniquely identifies each tuple:  $\forall t_1, t_2 \in R$  :  $(t_1[PK] = t_2[PK]) \Rightarrow (t_1 = t_2)$ . A foreign key  $FK_{R_2}$  in relation  $R_1$  then refers to the primary key PK of another relation  $R_2$  as  $\forall t \in R_1$  :  $t[FK] \in \{t'[PK] \mid t' \in R_2\}$ . This constitutes the interrelations in the database, with the RDBMS handling the referential integrity of  $T_{R_1}[FK] \subseteq T_{R_2}[PK]$ .

#### 2.2 Graph Neural Networks

Graph Neural Networks constitute a comprehensive class of neural models designed to process graph-structured data through the concept of (differentiable) *message-passing* [42]. Given an input graph  $G = (\mathcal{V}, \mathcal{E})$ , with a set of nodes  $\mathcal{V}$ and edges  $\mathcal{E}$ , let  $h_v^{(l)} \in \mathbb{R}^{d^{(l)}}$  be the vector representation (embedding) of node v at layer l. The general concept of GNNs can then be defined through the following sequence of three functions:

- 4 J. Peleška and G. Šír
- (i) Message function  $M^{(l)} : \mathbb{R}^{d^{(l)}} \times \mathbb{R}^{d^{(l)}} \to \mathbb{R}^{d_m^{(l)}}$  computes messages for each edge  $(u, v) \in E$  as  $m_{u \to v}^{(l)} = M^{(l)}(h_u^{(l)}, h_v^{(l)})$ .
- (ii) Aggregation function  $A^{(l)} : \{\mathbb{R}^{d_m^{(l)}}\} \to \mathbb{R}^{d_m^{(l)}}$  aggregates the messages for each  $v \in V$  as  $M_v^{(l)} = A^{(l)} \left( \{m_{u \to v}^{(l)} \mid (u, v) \in E\} \right)$ .
- (iii) Update function  $U^{(l)}: \mathbb{R}^{d^{(l)}} \times \mathbb{R}^{d_m^{(l)}} \to \mathbb{R}^{d^{(l+1)}}$  updates representation of each  $v \in V$  as  $h_v^{(l+1)} = U^{(l)}(h_v^{(l)}, M_v^{(l)})$ .

The specific choice of message, aggregation, and update functions varies across specific GNN models, which are typically structured with a predefined number L of such layers, enabling the message-passing to propagate information across L-neighborhoods within the graph(s).

### 2.3 Relational Deep Learning

In this paper, we adopt the concept of RDL as extending mainstream deep learning models, particularly the GNNs (Sec. 2.2), for application to RDBs (Sec. 2.1). For completeness, in the relational learning community [12], a number of similar approaches combining relational (logic-based) and deep learning methods arose under a similar name of "deep relational learning" [35]. Nevertheless, for compatibility with the recently introduced frameworks [15], we hereby continue with the contemporary RDL view, where RDBs are first transformed into a graph-based representation suitable for the GNN-based learning.

**Database Representation** The fundamental characteristic of RDL [15] is to represent an RDB as a heterogeneous graph.<sup>2</sup> The graph representation can be defined as  $G = (\mathcal{V}, \mathcal{E}, \mathcal{T}^v, \mathcal{T}^e)$ , where  $\mathcal{V}$  is the set of nodes,  $\mathcal{E}$  is the set of edges,  $\mathcal{T}^v$  is a set of node types with a mapping  $\phi : \mathcal{V} \to \mathcal{T}^v$ , and  $\mathcal{T}^e$  is a set of edge types with a mapping  $\psi : \mathcal{E} \to \mathcal{T}^e$ . The node types and edge types collectively form the graph *schema*  $(\mathcal{T}^v, \mathcal{T}^e)$ .

Given an RDB schema  $\mathcal{R}$ , the node types  $T \in \mathcal{T}^v$  correspond to the relations (tables) T within the database  $\mathcal{T}^v \xrightarrow{1:1} \mathcal{R}$ , while the edge types  $\mathcal{T}^e$  represent the undirected inter-relations between the tables, as defined by the primary-foreign key pairs:  $\mathcal{T}^e = \{(R_i, R_j) \mid R_i[FK_{R_j}] \subseteq R_j[PK] \lor R_j[FK_{R_i}] \subseteq R_i[PK]\}$ . For a specific *instance* of an RDB  $\mathcal{R}$ , the set of nodes  $\mathcal{V}$  is then defined as the union of all tuples (rows)  $t_i$  from each relation  $\mathcal{V} = \{v_{i,j} \mid R_i \in \mathcal{R}, t_j \in R_i\}$ , and the set of edges  $\mathcal{E}$  is defined as  $\mathcal{E} = \{(v_{i,k}, v_{j,l}) \mid t_k \in R_i, t_l \in R_j, (R_i, R_j) \in \mathcal{T}^e\}$ .

The graph representation is further enriched by node embedding matrices, attribute schema, and optionally a time mapping. Node embedding matrix  $h_v^{(l)} \in \mathbb{R}^{d \times d_{\phi(v)}}$  contains the embedding representation of a node  $v \in \mathcal{V}$  in a given layer *l*. With an attribute schema  $\mathcal{A}_T$  that provides information about the types of attributes  $A_1, \ldots, A_n$  associated with the nodes *v* of a specific node type  $T \in \mathcal{T}^v$ , the initial embedding tensors  $h_v^{(0)} \in \mathbb{R}^{d^{(0)} \times n}$  are computed from the

<sup>&</sup>lt;sup>2</sup> sometimes referred to as the "relational entity graph"

raw database attribute tuples  $t_i = (a_1, a_2, \ldots, a_n)$  through multi-modal attribute encoders [15]. Finally, the *time mapping* is a function  $\tau$  that assigns a timestamp  $t_v$  to each node  $\tau : v \mapsto t_v$ , effectively creating a dynamically growing graph in time, enabling the use of temporal graph sampling [34].

**Predictive Tasks** In RDL, predictive tasks are implemented through the creation of dedicated training tables  $T_t$  that extend the existing relational schema of  $\mathcal{R}$ . As introduced in [15], a training table  $T_t$  contains two essential components: foreign keys  $T_t[FK]$  that identify the entities of interest and target labels  $y \in \mathcal{A}_{T_t} \setminus T_t[FK]$ . Additionally, timestamps  $t_v \in \mathcal{A}_{T_t}$  that define temporal boundaries for the prediction of y can also be included.

The training table methodology supports a diverse range of predictive tasks, including node-level predictions (e.g., customer churn, product sales), link predictions between entities (e.g., user-product interactions), and, crucially, both temporal and static predictions. In the case of temporal predictions, a timestamp attribute  $t_v$  in the training table  $T_t$  specifies when the prediction is to be made, restricting the model to only consider information available up to the point  $t_v$  in time.

Neural Architecture Space Building upon the heterogeneous graph representation G, RDL models generally consist of the following four major stages.

- 1. Table-level attribute encoder creates the initial node embedding matrices  $h_v^{(0)} \in \mathbb{R}^{d^{(0)} \times n}$ , i.e. sequences of n embedding vectors  $\mathbb{R}^{d^{(0)}_{\phi(v)}}$  for each attribute  $A_1, \ldots, A_n$  of  $\phi(v)$  based on its respective semantic data type.
- 2. Table-level tabular model allows to employ existing tabular learning models [6,21] to yield more sophisticated node embeddings  $h_v^{(l)}$ . Notably, in this stage, an RDL model may reduce the dimensionality of the node at tribute matrix embedding  $h_v^{(l)} \in \mathbb{R}^{d^{(l)} \times n}$  to a vector embedding  $h_v^{(l)} \in \mathbb{R}^{d^{(l)}_{\phi(v)}}$ .
- 3. Graph neural model then depends on the chosen embedding dimensionality of  $h_v^{(l)}$ . If there is a single embedding vector  $h_v^{(l)} \in \mathbb{R}^{d_{\phi(v)}^{(l)}}$  per each node, the model can employ standard GNN (Sec. 2.2) heterogeneous messagepassing [38,4], otherwise a custom message-passing scheme [32] is required.
- 4. **Task-specific model head** finally provides transformation of the resulting node embeddings into prediction, usually involving simple MLP layers.

#### 3 The REDELEX Framework

The Relational Deep Learning Exploration (ReDeLEx) framework, which we introduce in this paper, offers a comprehensive environment for evaluating various RDL architectures over diverse RDB datasets.



Fig. 1: REDELEX end-to-end workflow for RDL.

### 3.1 Workflow Components

The ReDeLEx workflow, depicted in Fig. 1, consists of modular blocks that enable systematic exploration of the neural architecture and database configuration space, significantly extending the current scope [33] of RDL experimentation.

**Database Connectivity** In contrast to [33], the framework provides a standardized interface for connecting *directly* to an RDB [2], supporting various dialects of RDBMS. Notably, this enables a truly end-to-end deep learning pipeline, connecting to a possibly *remote* RDBMS, hosting the target RDB.

Attribute Schema Attribute schema creation, which is a crucial yet overlooked step in RDL, mediates information regarding the attribute types  $A_T$  within the specific node type  $T \in \mathcal{T}^v$  based on the original table attributes  $A \in \mathcal{A}_T$ . REDELEX automatically generates the attribute schema based on the SQL types and data from the RDBMS. Note that assessing a semantic type  $\mathsf{dom}(A_i)$  is not straightforward since, e.g., a SQL VARCHAR attribute  $A_i$  often stores categorical, textual, as well as temporal values  $a_i$ . To disambiguate such cases, we employ inbuilt heuristics utilizing the SQL types, names of the attributes, ratio of unique values, and patterns in the data to facilitate proper attribute embedding. **Predictive Tasks** The existing benchmark [33] provides support solely for tasks with a training table  $T_t$  generated from historical data through an SQL query. While useful, without any changes to the underlying database, this setting renders many RDB prediction tasks infeasible. REDELEX addresses this problem by adding support for tasks that require more substantial modifications of the original database. Tasks leveraging this functionality then not only generate a new table  $T_t$  but a whole modified instance  $\mathcal{R}'$  of the original database  $\mathcal{R}$ .

For example, assume the most common case where the database  $\mathcal{R}$  already contains the target attribute  $A_T$ , used for some node-level prediction task. In such a case the table  $T_t$  containing the target needs to be split into two tables  $T_{t_1}, T_{t_2}$  where  $T_{t_1}$  contains all original data except the target attribute  $A_T$  and is part of the newly modified database  $\mathcal{R}'$ , and  $T_{t_2}$  contains a duplicate of the primary key  $T_t[PK]$ , now used as a foreign key to the original table  $T_t$ , and the target attribute  $A_T$ . This table  $T_{t_2}$  is then used as the new training table  $T'_t$ .

Importantly, this scheme can be applied to generate tasks for unsupervised pretraining. Pretraining tasks can be created by choosing any table  $T \in \mathcal{R}$  in the database and duplicating it as T'. The unchanged duplicate T' can then be used as a training table  $T_t$ , while the values of cells in the original table T are randomly removed (masked out). The task is then to reconstruct any missing values in the classical tabular learning fashion [1], opening possibilities for sophisticated pretraining methods [36].

#### 3.2 RDL-Suitable Databases and Tasks

Due to the generality of the relational model, RDBs often contain data with vastly diverse structural characteristics that, in some cases, do not properly exploit the relational model (Sec. 2.1). Likewise, not all of the 70+ available RDBs [29] are actually suitable for the relational learning models. In this section, we examine RDB characteristics in the context of RDL to identify suitable databases to be used in the experiments (Sec. 4).

**Database Characteristics** To assess their overall characteristics, REDELEX associates each database task with various features pertinent to different parts if the training workflow (Fig. 1), which can be split into the below categories.

- 1. Database features provide high-level view of the data, including a domain (e.g. medicine, government, sport), whether the database is artificial or not, number of tables inside the database, number of foreign keys, number of factual (non-key) columns, number of columns with a specific variable type (e.g. numerical, categorical, time), total number of rows and total number of primary-foreign key pairs.
- 2. Schema features describe high-level structural aspects of the data. This includes the multiplicity of the relationships between the tables (one-to-one, one-to-many, many-to-many), features of the undirected graph induced by the primary-foreign key pairs (e.g. graph diameter,<sup>3</sup> or cycle detection).

<sup>&</sup>lt;sup>3</sup> Graph diameter is the maximum length of all the shortest paths between the nodes.

- 8 J. Peleška and G. Šír
- 3. Task features provide a similar type of information as the database features that are specific to the task and its target entity tables. This includes whether the task is temporal or static, number of training samples, multiplicity of relationships of the target entity table, etc.
- 4. Graph features inform about the properties of the transformed heterogeneous graph including, e.g., average eccentricity<sup>4</sup> of nodes or graph density.

**Tabular Data** A salient feature of RDBs are the inter-relations between the tables (Sec. 2.1). As such, it is obvious that RDBs that contain a single table, or multiple tables without any primary-foreign key pairs, will not benefit from the use of RDL. Furthermore, databases consisting of multiple tables linked solely by one-to-one relationships fall under the same category, as they allow for a complete *join* of the whole RDB into a single table. Importantly, as all values of foreign keys are unique (with the exception of missing values), all the resulting rows remain independent of each other, turning the RDL setting into standard tabular learning (see App. Tab. 4 for a list of such databases).

**Graph Data** On the other hand, RDBs are also characterized by building on the tabular representation, where an arbitrary number of attributes can be connected by a single relation. This is in contrast to the graph data which correspond to *binary* relational structures. Consequently, natively graph-structured data, such as molecules or family trees, although possible to be stored in an RDB, also do not fully exploit the relational model. In such cases, the RDL paradigm reduces to the simpler GNN setting [19], introducing an unnecessary complexity otherwise. More generally, RDL models for tasks on RDBs with a low number of non-key attributes (see Sec. 4.3) may suffer from information sparsity (see App. Table 5 for databases with the stated characteristics.)

# 4 Experiments

The aim of the experiments presented in this section is to demonstrate RE-DELEX in exploring the following selected RDL research questions:

- Q1: How do RDL methods perform in comparison to the traditional methods over diverse benchmarking tasks (Sec. 4.1)?
- Q2: Is it possible to apply tabular learning to a non-trivial RDB task while achieving results comparable to the RDL methods (Sec. 4.2)?
- Q3: What are some of the essential RDB characteristics that contribute to a successful application of a given learning model (Sec. 4.3)?

<sup>&</sup>lt;sup>4</sup> The eccentricity of a node is the maximum distance from the node to all other nodes.

**Databases** To establish a comprehensive yet manageable list from the overall 70+ available RDBs [29] for the RDL experimentation, we separated databases that exhibit the tabular (Sec. 3.2) or graph-like (Sec. 3.2) characteristics, or are artificially<sup>5</sup> created (see App. Table 6 for the most suitable databases).

**RDL Models** REDELEX is designed to accommodate development of highly diverse RDL architectures. For comprehensibility of the experiments, we present three models of gradually increasing complexity, selected from recent works. All the models fit into the outlined neural architecture space (Sec. 2.3), while utilizing the same attribute encoders for the numerical, categorical, multi-categorical, textual, and temporal values.

- 1. **GraphSAGE with Linear Transformation** is the simplest of the RDL models, applying a linear transformation on top of a concatenation of the attribute  $a_1, \ldots, a_n$  embeddings  $h_v^{(0)} \in \mathbb{R}^{n \cdot d^{(0)}}$  to yield a single embedding vector  $h_v^{(1)} = W h_v^{(0)}$  for each node v. The projected node embeddings  $h^{(1)} \in \mathbb{R}^{d_{\phi(v)}}$  then form input into the GraphSAGE [18] model, forming the GNN stage. Finally, a task specific model head is applied.
- 2. GraphSAGE with Tabular ResNet is similar to the previous, with the tabular-level stage reducing the node embedding dimensionality, however, the operation is performed through a more sophisticated tabular ResNet model [17]. Notably, this model was previously used in [33], allowing to directly align results between the RELBENCH and REDELEX benchmarks.
- 3. **DBFORMER** is an implementation of the Transformer-based RDL model from [32]. In constrast to the previous,<sup>6</sup> the model retains the original node embedding dimensionality  $h_v^{(l)} \in \mathbb{R}^{n \times d^{(l)}}$  while exploiting the attention mechanism [37] for learning interactions between both the attributes and tuples through a custom message-passing scheme.

Classical Models In addition to the selected RDL models, we include key representatives from related ML domains, including Gradient Boosted Decision Trees (GBDT; [31]), Deep Tabular Learning (DTL; [3]) and Propositionalization (Prop.; [24]). Particularly, we compare against the LightGBM [22] representative of GBDT; the getML's [11] FastProp feature generator combined with XGBoost [8]—representative of propositionalization; and the standalone tabular ResNet [17]—representative of deep tabular learning. Importantly, the LightGBM and the ResNet have access only to data from the task's target table, as these models fall into the tabular learning category. In contrast, the propositionalization method of FastProp with XGBoost exploits the full RDB structure.

 $<sup>^{5}</sup>$  with a single exception of the tpcd database

<sup>&</sup>lt;sup>6</sup> The key difference can be viewed analogously to the "fusion" and "cooperation" in multi-modal learning [20,27], considering the attributes as modalities. While the first two models *fuse* the representations of the discrete attributes at the beginning, the DBFORMER allows for *cooperation* of the attributes through the GNN stage.

|                       |        | GBDT  |                       | Prop. Tab |                       | Tab.   | Tab. DL |      | RDL                   |      |                       |         |                       |
|-----------------------|--------|-------|-----------------------|-----------|-----------------------|--------|---------|------|-----------------------|------|-----------------------|---------|-----------------------|
|                       | Model  | Light | LightGRM              |           | GetML                 |        | Tabular |      | ear                   | Res  | Net                   | DBEC    | DMED                  |
|                       | Widdei | Light | GDM                   | XGE       | $\mathbf{soost}$      | ResNet |         | SAGE |                       | SAGE |                       | DBIORME |                       |
| Database              | Task   | val   | $\operatorname{test}$ | val       | $\operatorname{test}$ | val    | test    | val  | $\operatorname{test}$ | val  | $\operatorname{test}$ | val     | $\operatorname{test}$ |
|                       |        |       |                       | Bin       | ary C                 | lassif | icatio  | n    |                       |      |                       |         |                       |
| ergastf1              | orig.  | .579  | .605                  | .086      | .081                  | .543   | .590    | .910 | .904                  | .913 | .907                  | .914    | .910                  |
| expend.               | orig.  | .852  | .856                  | .295      | .294                  | .846   | .847    | .920 | .918                  | .918 | .915                  | .894    | .892                  |
| geneea                | orig.  | .994  | .993                  | .317      | .258                  | .990   | .983    | .985 | .976                  | .986 | .981                  | .996    | .995                  |
|                       | temp.  | .997  | .990                  | .438      | .455                  | .989   | .982    | .984 | .971                  | .989 | .973                  | .998    | .991                  |
| hepatitis             | orig.  | .666  | .626                  | .044      | .055                  | .716   | .634    | 1.0  | .997                  | 1.0  | 1.0                   | 1.0     | .996                  |
| $\operatorname{imdb}$ | orig.  | .986  | .986                  | .451      | .454                  | .985   | .985    | .993 | .993                  | .993 | .993                  | .993    | .993                  |
| mondial               | orig.  | .500  | .500                  | NaN       | NaN                   | .500   | .542    | .987 | .988                  | 1.0  | .921                  | .987    | .932                  |
| movielens             | orig.  | .583  | .611                  | .223      | .211                  | .563   | .609    | .790 | .798                  | .801 | .809                  | .778    | .813                  |
| ${ m musk}$ L         | orig.  | .500  | .500                  | .381      | .240                  | .643   | .700    | .905 | .587                  | .937 | .547                  | .984    | .720                  |
| musk S                | orig.  | .500  | .500                  | .200      | .222                  | .625   | .583    | .867 | .889                  | .933 | .889                  | .800    | .759                  |
| mutagen.              | orig.  | .917  | .821                  | .095      | .024                  | .917   | .940    | .952 | .806                  | .964 | .798                  | .960    | .762                  |
| ncaa                  | orig.  | .687  | .426                  | .335      | .212                  | .692   | .559    | .848 | .759                  | .897 | .801                  | .826    | .720                  |
| stud.loan             | orig.  | .500  | .500                  | .272      | .247                  | .506   | .532    | 1.0  | 1.0                   | 1.0  | 1.0                   | 1.0     | 1.0                   |
| amazon                | ichurn | .738  | .740                  | NaN       | NaN                   | .653   | .654    | .823 | .826                  | .823 | .827                  | .818    | .821                  |
|                       | uchurn | .579  | .579                  | NaN       | NaN                   | .526   | .531    | .707 | .707                  | .707 | .707                  | .703    | .704                  |
| avito                 | clicks | .565  | .539                  | NaN       | NaN                   | .540   | .541    | .653 | .635                  | .654 | .672                  | .654    | .649                  |
|                       | visits | .536  | .528                  | NaN       | NaN                   | .520   | .516    | .702 | .662                  | .702 | .665                  | .699    | .660                  |
| f1                    | dnf    | .684  | .685                  | .324      | .230                  | .683   | .708    | .723 | .757                  | .744 | .740                  | .733    | .759                  |
|                       | top3   | .729  | .728                  | .300      | .158                  | .661   | .733    | .792 | .844                  | .714 | .832                  | .766    | .835                  |
| stack                 | badge  | .785  | .761                  | .083      | .097                  | .689   | .676    | .901 | .890                  | .901 | .891                  | .897    | .885                  |
|                       | engm.  | .842  | .832                  | .070      | .071                  | .788   | .782    | .902 | .906                  | .903 | .907                  | .902    | .903                  |
| trial                 | sout   | .676  | .711                  | .032      | .034                  | .624   | .664    | .662 | .702                  | .655 | .692                  | .645    | .675                  |
|                       |        |       |                       | Mult      | iclass                | Class  | ificat  | ion  |                       |      |                       |         |                       |
| accidents             | orig.  | .379  | .370                  | .506      | .505                  | .363   | .359    | 1.0  | 1.0                   | 1.0  | 1.0                   | 1.0     | 1.0                   |
|                       | temp.  | .186  | .170                  | .325      | .336                  | .204   | .187    | .819 | .583                  | .805 | .566                  | .825    | .727                  |
| craftbeer             | orig.  | .411  | .266                  | .003      | .017                  | .505   | .231    | .381 | .255                  | .488 | .329                  | .513    | .356                  |
| dallas                | orig.  | .594  | .581                  | NaN       | NaN                   | .585   | .572    | .547 | .499                  | .561 | .473                  | .688    | .511                  |
|                       | temp.  | .570  | .584                  | .288      | .512                  | .680   | .247    | .589 | .393                  | .615 | .424                  | .580    | .513                  |
| diabetes              | orig.  | .190  | .190                  | .402      | .383                  | .190   | .190    | .878 | .870                  | .892 | .883                  | .890    | .875                  |
| financial             | orig.  | .461  | .458                  | .899      | .658                  | .698   | .470    | .586 | .490                  | .738 | .729                  | .644    | .539                  |
|                       | temp.  | .495  | .492                  | .636      | .632                  | .456   | .449    | .495 | .492                  | .538 | .340                  | .553    | .256                  |
| genes                 | orig.  | .060  | .060                  | .066      | .100                  | .060   | .060    | 1.0  | .898                  | 1.0  | 1.0                   | 1.0     | .927                  |
| hockey                | orig.  | .707  | .721                  | .395      | .404                  | .707   | .703    | .714 | .728                  | .714 | .724                  | .713    | .740                  |
| legalacts             | orig.  | .933  | .924                  | .320      | .316                  | .837   | .824    | .834 | .822                  | .837 | .825                  | .831    | .821                  |
|                       | temp.  | .870  | .851                  | .226      | .220                  | .788   | .688    | .791 | .721                  | .786 | .698                  | .784    | .703                  |
| premiere-             | orig.  | .380  | .411                  | .641      | .601                  | .589   | .362    | .659 | .400                  | .593 | .332                  | .571    | .352                  |
| league                | temp.  | .317  | .222                  | .647      | .639                  | .615   | .370    | .650 | .384                  | .619 | .390                  | .556    | .407                  |
| tpcd                  | orig.  | .191  | .182                  | NaN       | NaN                   | .191   | .185    | .733 | .732                  | .546 | .545                  | .722    | .728                  |
| webkp                 | orig.  | .139  | .130                  | .350      | .366                  | .139   | .130    | .315 | .193                  | .295 | .212                  | .355    | .282                  |
|                       |        | 25    | 38                    | 1         | 21                    | 19     | 18      | 26   | 37                    | 2    | 51                    | 2       | 17                    |
| Avg. Rank             |        | 2.4   | 43                    | 3.3       | 28                    | 2      | 8       | 2.0  |                       | 1.   | 36                    |         |                       |
|                       |        | · -·· |                       |           |                       |        | ~       | 1    |                       |      |                       |         |                       |

Table 1: Overall results from the classification tasks, presenting AUC ROC values for the binary classification, and macro f1 score for the multiclass classification, respectively (higher is better). Static (non-temporal) tasks are tagged as "orig."

#### 4.1 Benchmarking tasks

We present comprehensive results over two types of node-level tasks—binary classification and multiclass classification. The tasks can be further differentiated by the origin of the target labels and usage of temporal values. Tasks performed on datasets from the RELBENCH collection use generated target table attributes, while tasks on datasets from the CTU Relational use existing target table attributes. Additionally, tasks from the CTU Relational can be both static and temporal, while tasks from the RELBENCH collection are always temporal. Static and temporal tasks differentiate based on the constrains forced upon the sampling algorithm while generating a sub-subgraph used for training the model. and by the method of splitting the dataset between the training, validation and test data. Static tasks use neighborhood sampling [18] constrained only by the maximum number of neighbors, and data splitting is carried out at random w.r.t. a given ratio (e.g. 70:15:15). In contrast, temporal tasks extend the neighborhood sampling by incorporating temporal constraints, assuming only directed edges from nodes with an older timestamp and, similarly, the splits are carried out w.r.t. the timestamps where all training entities must precede validation and testing data, respectively (see App. A for the full experimental setup).

The results on classification tasks shown in Table 1 demonstrate a strong performance of the RDL models (Sec. 4) over the classical models (Sec. 4) on majority of the datasets. Specifically, the ResNet SAGE performs very well on binary classification tasks, especially on tasks with a training table generated from historical data. The higher complexity of the DBFORMER somewhat surprisingly seems to lead to higher robustness as the model slightly outperforms the ResNet SAGE in the average rank. Nevertheless, in a couple of cases, the simplest Linear SAGE model matches or even outperforms the score of both the models. Notably, on the hepatits, mondial, student loan, accidents, and genes datasets, all the RDL models present near-perfect predictions while the classical models show an order of magnitude worse score, highlighting the contribution of the RDL representation.

#### 4.2 Tabular Learning

In this scenario, we compare results of the Tabular Learning (TL) models from the previous section to ones trained on new tables formed by join operations over the RDBs (Sec. 3.2). Particularly, we evaluate the TL models on tables generated by joins over *both* one-to-one and many-to-one relationships of the target table.<sup>7</sup> Additionally, we include evaluation of RDL models with exactly 2 layers in the graph neural stage, which is conceptually equivalent to the join operation. Finally, we include the overall best RDL models to put the results into context. Note that the previous Table 1 demonstrated that the TL methods perform significantly worse on an absolute majority of tasks. In this experiment, we aim to assess whether a simple RDB transformation could actually change

<sup>&</sup>lt;sup>7</sup> This is similar to a recent RelGNN method [7], albeit limited to the target table.

|                              | Model                   | Ligh |                       |            |                       | ResNet Tabula |                       |            | ır                    |          | RD                    | DL   |      |
|------------------------------|-------------------------|------|-----------------------|------------|-----------------------|---------------|-----------------------|------------|-----------------------|----------|-----------------------|------|------|
|                              | model                   | Base |                       | Joined     |                       | Base          |                       | Joined     |                       | 2 layers |                       | Best |      |
| Database                     | Task                    | val  | $\operatorname{test}$ | val        | $\operatorname{test}$ | val           | $\operatorname{test}$ | val        | $\operatorname{test}$ | val      | $\operatorname{test}$ | val  | test |
|                              |                         |      |                       | Binar      | ry Cla                | assific       | ation                 |            |                       |          |                       |      |      |
| ergastf1                     | orig.                   | .579 | .605                  | .919       | .904                  | .543          | .590                  | .899       | .892                  | .909     | .905                  | .910 | .904 |
| hepatitis                    | orig.                   | .666 | .626                  | .940       | .913                  | .716          | .634                  | .718       | .632                  | 1.0      | 1.0                   | 1.0  | 1.0  |
| mondial                      | orig.                   | .500 | .500                  | <u>1.0</u> | .815                  | .500          | .542                  | <u>1.0</u> | .898                  | .988     | .944                  | .988 | .944 |
| movielens                    | orig.                   | .583 | .611                  | .574       | .622                  | .563          | .609                  | .563       | .609                  | .688     | .691                  | .784 | .796 |
| ncaa                         | orig.                   | .687 | .426                  | .797       | .776                  | .692          | .559                  | .819       | .741                  | .801     | .728                  | .861 | .764 |
| $\operatorname{studentloan}$ | orig.                   | .500 | .500                  | <u>1.0</u> | <u>1.0</u>            | .506          | .532                  | <u>1.0</u> | <u>1.0</u>            | 1.0      | 1.0                   | 1.0  | 1.0  |
| amazon                       | $\operatorname{uchurn}$ | .579 | .579                  | .577       | .577                  | .526          | .531                  | .526       | .531                  | .705     | .705                  | .706 | .707 |
| avito                        | clicks                  | .565 | .539                  | .564       | .556                  | .540          | .541                  | .540       | .541                  | .647     | .658                  | .653 | .649 |
|                              | visits                  | .536 | .528                  | .536       | .532                  | .520          | .516                  | .520       | .516                  | .690     | .656                  | .702 | .664 |
| stack                        | badge                   | .785 | .761                  | .783       | .759                  | .689          | .676                  | .689       | .676                  | .898     | .887                  | .901 | .891 |
|                              |                         |      | Ν                     | Iultic     | lass (                | Classif       | ficatio               | n          |                       |          |                       |      |      |
| accidents                    | orig.                   | .379 | .370                  | .377       | .372                  | .363          | .359                  | .366       | .365                  | 1.0      | 1.0                   | 1.0  | 1.0  |
|                              | temp.                   | .186 | .170                  | .177       | .163                  | .204          | .187                  | .198       | .178                  | .790     | .579                  | .798 | .582 |
| $\operatorname{craftbeer}$   | orig.                   | .411 | .266                  | .488       | .245                  | .505          | .231                  | .505       | .231                  | .442     | .301                  | .444 | .311 |
| diabetes                     | orig.                   | .190 | .190                  | .190       | .190                  | .190          | .190                  | .190       | .190                  | .885     | .876                  | .885 | .876 |
| genes                        | orig.                   | .060 | .060                  | .060       | .060                  | .060          | .060                  | .060       | .060                  | .989     | .922                  | 1.0  | .896 |
| premiere-                    | orig.                   | .380 | .411                  | .470       | .416                  | .589          | .362                  | .584       | .325                  | .603     | .371                  | .603 | .371 |
| league                       | temp.                   | .317 | .222                  | .527       | .485                  | .615          | .370                  | .598       | .399                  | .616     | .431                  | .616 | .431 |
| tpcd                         | orig.                   | .191 | .182                  | .846       | <u>.841</u>           | .191          | .185                  | .676       | .666                  | .579     | .582                  | .579 | .582 |
| webkp                        | orig.                   | .139 | .130                  | .139       | .130                  | .139          | .130                  | .139       | .130                  | .265     | .185                  | .308 | .221 |

Table 2: Classification tasks over a subset of datasets formed by joining the target table, showing AUC ROC values for binary classification, and macro f1 score for multiclass classification, respectively (higher is better). Significant improvements (more than 0.05 score) are shown in bold, while new best results are underlined.

the situation in some cases. Particularly, we select a subset of tasks from Table 1 where both the TL models showed at least 0.1 worse score than the best RDL model. The results in Table 2 show that, indeed, on a number of datasets the TL models register a significant improvement with results sometimes comparable to the best of RDL. Notably, on the ncaa, premierleague and tpcd they even set new best results. This experiment demonstrates existing weak spots in the new RDL approach [15], suggesting that caution and thorough analysis are still in order before deploying RDL on an RDB task.

#### 4.3 Essential Characteristics

The overarching aim of REDELEX is to assess common characteristics of RDBs and tasks in the context of a used learning approach. While a fully comprehensive assessment is out of scope of this short paper, in Table 3 we summarize characteristics of the databases against the respective performances of the vari-

| Fosture    | Base  |       | RDL    |                   | Pro   | position | ıal.              | Tabular Learning |        |                   |
|------------|-------|-------|--------|-------------------|-------|----------|-------------------|------------------|--------|-------------------|
| reature    | Med.  | Q1    | Mediar | 1 Q3              | Q1    | Mediar   | 1 Q3              | Q1               | Median | Q3                |
| #Tab.      | 6     | 3     | 7      | 8                 | 3     | 4        | 8                 | 3                | 4      | 13.5              |
| #FK        | 6     | 2.25  | 6      | 12                | 2.5   | 5        | 8                 | 3                | 4      | 14                |
| #Factual   | 35    | 14    | 24     | 45                | 17    | 47       | 110               | 27               | 40     | 119               |
| #Cat. ∥    | 11    | 2.25  | 11     | 16                | 6     | 17       | 44                | 9                | 14     | 59                |
| #Num. ∣    | 6     | 2     | 4.5    | 13                | 4     | 6        | 12                | 2.5              | 3      | 9                 |
| #Text      | 6     | 2     | 6      | 10                | 1     | 2        | 3.5               | 3                | 14     | 20                |
| #Time      | 2     | 0     | 2      | 7                 | 0.5   | 3        | 5                 | 1.5              | 3      | 5                 |
| #Rows      | 472k  | 96k   | 1.36M  | $5.59 \mathrm{M}$ | 16.6k | 81.8k    | 1.6M              | 32.1k            | 804k   | $2.54 \mathrm{M}$ |
| #Links     | 752k  | 182k  | 2.22M  | $7.98 \mathrm{M}$ | 32.5k | 82.6k    | $1.58 \mathrm{M}$ | 45.6             | 851k   | 3.49M             |
| Diameter   | 3     | 2     | 2.5    | 3                 | 2     | 2        | 3                 | 2                | 3      | 4.5               |
| 1-to-1     | 0     | 0     | 0      | 2                 | 0     | 0        | 4                 | 0                | 2      | 3                 |
| 1-to-M     | 5     | 2     | 4.5    | 11                | 2.5   | 5        | 5                 | 2.5              | 4      | 12.5              |
| #Train     | 12k   | 2.22k | 27.1k  | $1.15 \mathrm{M}$ | 305   | 546      | 21.8k             | 1.96k            | 18.5k  | 348k              |
| #T.Factual | 5     | 1     | 3      | 6                 | 2.5   | 5        | 5                 | 4                | 10     | 15                |
| #T.Cat     | 2     | 0     | 0      | 2                 | 0.5   | 1        | 4                 | 2                | 4      | 10                |
| #T.Num     | 0     | 0     | 0      | 1                 | 0     | 0        | 2                 | 0                | 1      | 2                 |
| #T.Text    | 1     | 0     | 1      | 3                 | 0     | 0        | 1.5               | 0                | 1      | 5.5               |
| #T.Time    | 1     | 0     | 0      | 1                 | 0     | 1        | 1                 | 1                | 1      | 2                 |
| Eccentric. | 10.6  | 7.13  | 15.3   | 19.5              | 3.3   | 9.7      | 15.8              | 3.39             | 10.6   | 17.4              |
| Density    | 0.025 | 0.001 | 0.018  | 0.079             | 0.004 | 0.028    | 0.102             | 0.003            | 0.059  | 0.118             |

Table 3: Characteristics of databases and their tasks selected based on the best performing model. Features are sorted into the categories described in Sec. 3.2.

ous model types. Following the analysis, RDL models generally tend to perform well on datasets with a large number of training samples and links. Propositionalization achieves best results mostly on smaller datasets with a low number of one-to-many relationships, yet with a large number of factual (non-key) columns. This is in line with some previous studies [26], despite the remaining prevalence of propositionalization methods in practice [11]. The TL methods then tend to perform best when there is a higher number of factual columns in the target table, which aligns with natural intuition. Moreover, these allow capturing more diverse attribute types where, e.g., both LightGBM and deep TL models are capable of utilizing textual attributes.

**Related Work** As outlined in the Introduction, the REDELEX framework builds upon the CTU relational dataset collection [29] which it integrates with the RELBENCH [33] interface to facilitate a wider scope of RDL [15] experimentation. As such, it is naturally related to recent works introducing new RDL models, which include [32,43,7]. Besides RDL, related work includes other dataset and benchmarking frameworks that address some facets of learning from relational data, including [39,40]. The most salient feature of REDELEX, within the context of related work, is the provided bridge between the traditional relational learning methods [14] and the contemporary RDL [15].

### 5 Conclusion

In this study, we introduced REDELEX—a framework for exploring and evaluating Relational Deep Learning (RDL) models across diverse relational database contexts. The framework enables benchmarking on more than 70 databases, facilitating new insights into the relationships between the RDL neural architecture choices, traditional learning methods, and the underlying database characteristics. Our results demonstrated that RDL approaches mostly outperform the traditional methods. Nevertheless, a closer inspection revealed important cases in which the performance of the competing tabular learning methods could be easily improved to match or even surpass RDL, highlighting the interim immaturity of the field, and the need for further RDL exploration. Our general exploration in this paper demonstrated that RDL performs well on databases with complex relationships and large numbers of samples, while the traditional methods may still remain a sensible choice for smaller and flatter datasets.

Ethical Considerations The performance of RDL models demonstrated in our research could enable more sophisticated inference of personal information from interconnected data sources. The framework's flexibility could facilitate deployment in domains with significant ethical implications, such as financial services, healthcare, and government operations. We encourage researchers using REDELEX to carefully assess privacy implications and implement appropriate anonymization techniques.

Acknowledgments. This work has received funding from the European Union's Horizon Europe Research and Innovation program under the grant agreement TUPLES No. 101070149; and Czech Science Foundation grant No. 24-11664S.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

### References

- Arik, S.Ö., Pfister, T.: Tabnet: Attentive interpretable tabular learning. In: Proceedings of the AAAI conference on artificial intelligence. pp. 6679–6687 (2021)
- 2. Bayer, M.: Sqlalchemy. In: The Architecture of Open Source Applications Volume II: Structure, Scale, and a Few More Fearless Hacks. aosabook.org (2012)
- Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., Kasneci, G.: Deep neural networks and tabular data: A survey. IEEE transactions on neural networks and learning systems (2022)
- 4. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? In: International Conference on Learning Representations (2022)
- Chamberlin, D.D., Boyce, R.F.: SEQUEL: A structured English query language. In: Proceedings of the 1974 ACM SIGFIDET (now SIGMOD) workshop on Data description, access and control. pp. 249–264. SIGFIDET '74, Association for Computing Machinery, New York, NY, USA (1974)
- 6. Chen, K.Y., Chiang, P.H., Chou, H.R., Chen, T.W., Chang, T.H.: Trompt: Towards a better deep neural network for tabular data (2023)

- Chen, T., Kanatsoulis, C., Leskovec, J.: RelGNN: Composite Message Passing for Relational Deep Learning. arXiv preprint arXiv:2502.06784 (2025)
- Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining. pp. 785–794 (2016)
- Codd, E.F.: A relational model of data for large shared data banks. Commun. ACM 13(6), 377–387 (1970)
- Codd, E.F.: The relational model for database management: version 2. Addison-Wesley Longman Publishing Co., Inc. (1990)
- 11. Code17 GmbH: getml, https://getml.com
- Cropper, A., Dumančić, S., Muggleton, S.H.: Turning 30: New ideas in inductive logic programming. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20. pp. 4833–4839 (2020)
- Cvitkovic, M.: Supervised learning on relational databases with graph neural networks (2020)
- 14. De Raedt, L.: Logical and Relational Learning. Springer (2008)
- Fey, M., Hu, W., Huang, K., Lenssen, J.E., Ranjan, R., Robinson, J., Ying, R., You, J., Leskovec, J.: Position: Relational deep learning - graph representation learning on relational databases. In: Forty-first International Conference on Machine Learning (2024)
- Gallier, J.H.: Logic for computer science: foundations of automatic theorem proving. Courier Dover Publications (2015)
- Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting deep learning models for tabular data. In: Proceedings of the 35th International Conference on Neural Information Processing Systems. pp. 18932–18943. NIPS '21, Curran Associates Inc., Red Hook, NY, USA (2021)
- Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: Advances in neural information processing systems. pp. 1024–1034 (2017)
- Hamilton, W.L.: Graph Representation Learning. Morgan & Claypool Publishers (2020)
- Hu, R., Singh, A.: Unit: Multimodal multitask learning with a unified transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1439–1449 (2021)
- Hu, X., Tang, W., Hsieh, C.K., Shi, S.: Tabtransformer: Tabular data modeling using contextual embeddings. arXiv preprint arXiv:2012.06678 (2020)
- 22. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: LightGBM: a highly efficient gradient boosting decision tree. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 3149–3157. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (2017)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Kramer, S., Lavrač, N., Flach, P.: Propositionalization approaches to relational data mining. Relational data mining pp. 262–291 (2001)
- Krogel, M.A., Rawles, S., Żelezný, F., Flach, P.A., Lavrač, N., Wrobel, S.: Comparative evaluation of approaches to propositionalization. Springer (2003)
- Lavrač, N., Podpečan, V., Robnik-Sikonja, M.: Machine Learning Background, pp. 17–53. Springer International Publishing, Cham (2021)
- Liang, P.P., Zadeh, A., Morency, L.P.: Foundations & trends in multimodal machine learning: Principles, challenges, and open questions. ACM Computing Surveys 56(10), 1–42 (2024)

- 16 J. Peleška and G. Šír
- Maali, F., Cyganiak, R., Peristeras, V.: Enabling Interoperability of Government Data Catalogues. In: Wimmer, M.A., Chappelet, J.L., Janssen, M., Scholl, H.J. (eds.) Electronic Government. pp. 339–350. Springer, Berlin, Heidelberg (2010)
- Motl, J., Schulte, O.: The ctu prague relational learning repository. arXiv preprint arXiv:1511.03086 (2015)
- Muggleton, S., De Raedt, L.: Inductive logic programming: Theory and methods. The Journal of Logic Programming 19 (1994)
- Natekin, A., Knoll, A.: Gradient boosting machines, a tutorial. Frontiers in neurorobotics 7, 21 (2013)
- Peleška, J., Šír, G.: Transformers meet relational databases. arXiv preprint arXiv:2412.05218 (2024)
- 33. Robinson, J., Ranjan, R., Hu, W., Huang, K., Han, J., Dobles, A., Fey, M., Lenssen, J.E., Yuan, Y., Zhang, Z., He, X., Leskovec, J.: Relbench: A benchmark for deep learning on relational databases. In: The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2024)
- Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., Bronstein, M.: Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637 (2020)
- Sír, G.: Deep Learning with Relational Logic Representations. Czech Technical University (2021)
- Somepalli, G., Goldblum, M., Schwarzschild, A., Bruss, C.B., Goldstein, T.: Saint: Improved neural networks for tabular data via row attention and contrastive pretraining (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems (2017)
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: International Conference on Learning Representations (2018)
- Vogel, L., Bodensohn, J.M., Binnig, C.: WikiDBs: A Large-Scale Corpus Of Relational Databases From Wikidata. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) Advances in Neural Information Processing Systems. vol. 37, pp. 41186–41201. Curran Associates, Inc. (2024)
- Wang, M., Gan, Q., Wipf, D., Zhang, Z., Faloutsos, C., Zhang, W., Zhang, M., Cai, Z., Li, J., Mao, Z., et al.: 4dbinfer: A 4d benchmarking toolbox for graph-centric predictive modeling on rdbs. Advances in Neural Information Processing Systems 37, 27236–27273 (2024)
- White, J.: PubMed 2.0. Medical Reference Services Quarterly 39(4), 382–387 (2020)
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Philip, S.Y.: A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems (2020)
- Zahradník, L., Neumann, J., Sír, G.: A deep learning blueprint for relational databases. In: NeurIPS 2023 Second Table Representation Learning Workshop (2023)
- 44. Zhang, H., Gan, Q., Wipf, D., Zhang, W.: Gfs: Graph-based feature synthesis for prediction over relational databases. arXiv preprint arXiv:2312.02037 (2023)

# A Experimental Setup

All deep learning models (including RDL) were trained for a minimum of 10 epochs with a total minimum of 1000 steps of Adam [23] optimizer with a learning rate of 0.001. Experiments on the RDL models were conducted with fixed hyperparameters with two exceptions—the number of layers of the graph neural model and the neighborhood graph sampling rate, which were searched for in a grid hyperparameter optimization. The hyperparameters include batch size—set to 512, message-passing aggregation function—set to summation, embedding vectors dimension, which is the same for both row and attribute embedding vectors—set to 64, neighborhood sampling rate—iterated over the values of 16, 32 and 64, and number of message-passing layers—a value in range of 1 to 4.

### **B** Additional tables

Here we provide additional tables with descriptive information about the databases available through REDELEX.

| Database        | Domain  | #Tables | # FK | #Factual | one-to-one | many-to-one |
|-----------------|---------|---------|------|----------|------------|-------------|
| atherosclerosis | Med.    | 4       | 3    | 191      | 2          | 1           |
| bupa            | Med.    | 9       | 8    | 16       | 7          | 1           |
| cde             | Gov.    | 3       | 2    | 87       | 2          | 0           |
| pima            | Med.    | 9       | 8    | 9        | 8          | 0           |
| satellite       | Indstr. | 34      | 34   | 67       | 34         | 0           |
| voc             | Retail  | 8       | 7    | 89       | 6          | 1           |

Table 4: Tabular-like databases available through REDELEX.

| Database       | Domain   | #Tables | # FK | #Factual | one-to-one | many-to-one |
|----------------|----------|---------|------|----------|------------|-------------|
| carcinogenesis | Med.     | 6       | 13   | 4        | 2          | 11          |
| cora           | Edu.     | 3       | 3    | 2        | 0          | 3           |
| mesh           | Industry | 29      | 33   | 37       | 24         | 9           |
| pima           | Med.     | 9       | 8    | 9        | 8          | 0           |
| toxicology     | Med.     | 4       | 5    | 3        | 0          | 5           |

Table 5: Graph-like databases available through REDELEX.

| Database                 | Dom.   | #Tab. | # FK   | #Fact.  | Diam. | Cycle | 1:1 | 1:N |  |  |  |
|--------------------------|--------|-------|--------|---------|-------|-------|-----|-----|--|--|--|
| CTU Relational Databases |        |       |        |         |       |       |     |     |  |  |  |
| accidents                | Gov.   | 3     | 3      | 38      | 1     | ~     | 0   | 3   |  |  |  |
| basketballmen            | Sport  | 9     | 9      | 187     | 5     | 1     | 2   | 8   |  |  |  |
| biodegradability         | Med.   | 5     | 5      | 6       | 3     | 1     | 0   | 5   |  |  |  |
| countries                | Geo.   | 3     | 2      | 63      | 2     | ×     | 0   | 2   |  |  |  |
| craftbeer                | Entmt. | 2     | 1      | 8       | 1     | ×     | 0   | 1   |  |  |  |
| dallas                   | Gov.   | 3     | 2      | 24      | 2     | ×     | 0   | 2   |  |  |  |
| diabetes                 | Edu.   | 3     | 2      | 4       | 2     | X     | 0   | 2   |  |  |  |
| ergastf1                 | Sport  | 13    | 19     | 82      | 3     | 1     | 0   | 19  |  |  |  |
| expenditures             | Retail | 3     | 2      | 19      | 2     | ×     | 0   | 2   |  |  |  |
| financial                | Fin.   | 8     | 8      | 39      | 3     | 1     | 6   | 5   |  |  |  |
| fnhk                     | Med.   | 3     | 2      | 21      | 2     | ×     | 0   | 2   |  |  |  |
| geneea                   | Gov.   | 19    | 20     | 99      | 7     | 1     | 8   | 16  |  |  |  |
| genes                    | Med.   | 3     | 3      | 11      | 2     | 1     | 0   | 3   |  |  |  |
| grants                   | Edu.   | 12    | 11     | 30      | 4     | X     | 2   | 10  |  |  |  |
| hepatitis                | Med.   | 7     | 6      | 16      | 4     | ×     | 6   | 3   |  |  |  |
| hockey                   | Sport  | 19    | 27     | 273     | 4     | 1     | 8   | 23  |  |  |  |
| imdb                     | Entmt. | 7     | 6      | 12      | 5     | X     | 0   | 6   |  |  |  |
| lahman                   | Sport  | 25    | 31     | 319     | 6     | 1     | 6   | 28  |  |  |  |
| legalacts                | Gov.   | 4     | 4      | 24      | 3     | 1     | 0   | 4   |  |  |  |
| mondial                  | Geo.   | 33    | 62     | 125     | 5     | 1     | 12  | 55  |  |  |  |
| movielens                | Entmt. | 7     | 6      | 14      | 4     | X     | 0   | 6   |  |  |  |
| musklarge                | Med.   | 2     | 1      | 167     | 1     | X     | 0   | 1   |  |  |  |
| musksmall                | Med.   | 2     | 1      | 167     | 1     | ×     | 0   | 1   |  |  |  |
| mutagenesis              | Med.   | 3     | 3      | 9       | 2     | 1     | 0   | 3   |  |  |  |
| ncaa                     | Sport  | 8     | 15     | 99      | 3     | 1     | 0   | 15  |  |  |  |
| premiereleague           | Sport  | 4     | 5      | 209     | 2     | 1     | 0   | 5   |  |  |  |
| restbase                 | Retail | 3     | 3      | 10      | 1     | 1     | 2   | 2   |  |  |  |
| seznam                   | Retail | 4     | 3      | 10      | 2     | ×     | 0   | 3   |  |  |  |
| sfscores                 | Gov.   | 3     | 2      | 22      | 2     | ×     | 0   | 2   |  |  |  |
| stats                    | Edu.   | 8     | 13     | 50      | 3     | 1     | 2   | 12  |  |  |  |
| studentloan              | Fin.   | 10    | 9      | 15      | 3     | ×     | 12  | 3   |  |  |  |
| tpcd                     | Retail | 8     | 8      | 48      | 4     | 1     | 1   | 7   |  |  |  |
| triazine                 | Med.   | 2     | 1      | 13      | 1     | ×     | 0   | 1   |  |  |  |
| walmart                  | Retail | 4     | 3      | 27      | 3     | ×     | 0   | 3   |  |  |  |
| webkp                    | Edu.   | 3     | 3      | 5       | 2     | 1     | 0   | 3   |  |  |  |
|                          |        | RelB  | ench L | atabase | es    |       |     |     |  |  |  |
| amazon                   | Retail | 3     | 2      | 10      | 2     | X     | 0   | 2   |  |  |  |
| avito                    | Retail | 8     | 11     | 23      | 2     | 1     | 0   | 11  |  |  |  |
| f1                       | Sport  | 9     | 13     | 45      | 3     | 1     | 0   | 13  |  |  |  |
| stack                    | Edu.   | 7     | 12     | 33      | 3     | 1     | 2   | 11  |  |  |  |
| trial                    | Med.   | 15    | 15     | 110     | 4     | ~     | 4   | 13  |  |  |  |

Table 6: List of databases available for benchmarking in REDELEX.