

WarpGate: A Semantic Join Discovery System for Cloud Data Warehouses

Tianji Cong*
University of Michigan
Ann Arbor, Michigan, USA
congjtj@umich.edu

James Gale
Sigma Computing
San Francisco, California, USA
jlg@sigmacomputing.com

Jason Frantz
Sigma Computing
San Francisco, California, USA
jason@sigmacomputing.com

H. V. Jagadish
University of Michigan
Ann Arbor, Michigan, USA
jag@umich.edu

Çağatay Demiralp
Sigma Computing
San Francisco, California, USA
cagatay@sigmacomputing.com

ABSTRACT

Data discovery is a major challenge in enterprise data analysis: users often struggle to find data relevant to their analysis goals or even to navigate through data across data sources, each of which may easily contain thousands of tables. One common user need is to discover tables joinable with a given table. This need is particularly critical because join is a ubiquitous operation in data analysis, and join paths are mostly obscure to users, especially across databases. Furthermore, users are typically interested in finding “semantically” joinable tables: with columns that can be transformed to become joinable even if they are not joinable as currently represented in the data store.

We present WARP_{GATE}, a system prototype for data discovery over cloud data warehouses. WARP_{GATE} implements an embedding-based solution to semantic join discovery, which encodes columns into high-dimensional vector space such that joinable columns map to points that are near each other. Through experiments on several table corpora, we show that WARP_{GATE} (i) captures semantic relationships between tables, especially those across databases, and (ii) is sample efficient and thus scalable to very large tables of millions of rows. We also showcase an application of WARP_{GATE} within an enterprise product for cloud data analytics.

1 INTRODUCTION

Three major correlated forces act on enterprise data analysis today. First, more and more enterprise data are stored on the cloud. Second, the size and the diversity of data available to enterprises increase faster than ever as a growing number of data sources feed into enterprises’ cloud data warehouses (CDWs). Third, the number of users in enterprises who would like direct access to cloud data for decision-making is also rapidly growing, boosting the demand for easy-to-use cloud data analysis systems such as Sigma Workbooks [15]. Hence enabling easy and effective data discovery is increasingly more valuable. However, existing commercial solutions such as data catalogs are typically inaccessible from business

intelligence tools, inadequate for surfacing relationships unknown to users, and inflexible in capturing data semantics.

Not surprisingly, recent years have seen a rising interest in data discovery, especially for Open Data [2, 19, 21, 25]. This is a challenging problem due to the unstructured/semi-structured nature of Open Data and the lack of high-quality metadata. However, we observe that modern organizations also face data discovery challenges when dealing with structured data (i.e., databases in CDWs) for similar reasons: (i) in large CDWs, few people have a good understanding of data and their relationships; (ii) data is added in a much less modeled state as cloud computing power and storage become easily accessible; and (iii) business users who do not have comprehensive knowledge of underlying schema designs, are usually unaware of relationships between datasets (e.g., PK/FK relationship between two tables). There is a need for surfacing such relationships to help users connect related datasets within and across databases. We also observe that CDW data tends to have more structure than Open Data, sometimes has high update rates, and usually has far more stringent requirements for completeness of discovery results.

Consider a business user Joey who has access to a CDW that comprises hundreds of databases and thousands of tables (Figure 1). In the interest of launching a sales campaign and finding potential customers, she starts with the ACCOUNT table in their SALESFORCE database, which contains member and billing data of their customer companies. To select ideal customers for the campaign, she wants to augment the ACCOUNT table with additional information such as sectors in which companies are engaged in their business activities. Now she has two options to find that complementary information: either she manually browses tables in the same database (even wanders through the entire data warehouse), or she can ask internal data teams for support. Either way, her workflow is significantly delayed. There is also a fair chance that she cannot go through all databases as it is a tedious and laborious process, and the data team does not have that particular piece of knowledge.

Despite efforts in prototyping data discovery for real-world applications [1, 3, 12], aspects of data discovery pertaining to enterprise settings are overlooked. For example, previous works assume a full pass of underlying data in the indexing/profiling phase. However, it is computationally and monetarily expensive to pull all data out of CDWs. In addition, existing solutions, including those for Open Data, focus on improving index lookup time for the scale of Open Data lakes or Web table corpora, which usually consist of a very

*Work done when Tianji was interning at Sigma Computing.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution, provided that you attribute the original work to the authors and CIDR 2023. 13th Annual Conference on Innovative Data Systems Research (CIDR '23), January 8-11, 2023, Amsterdam, The Netherlands.

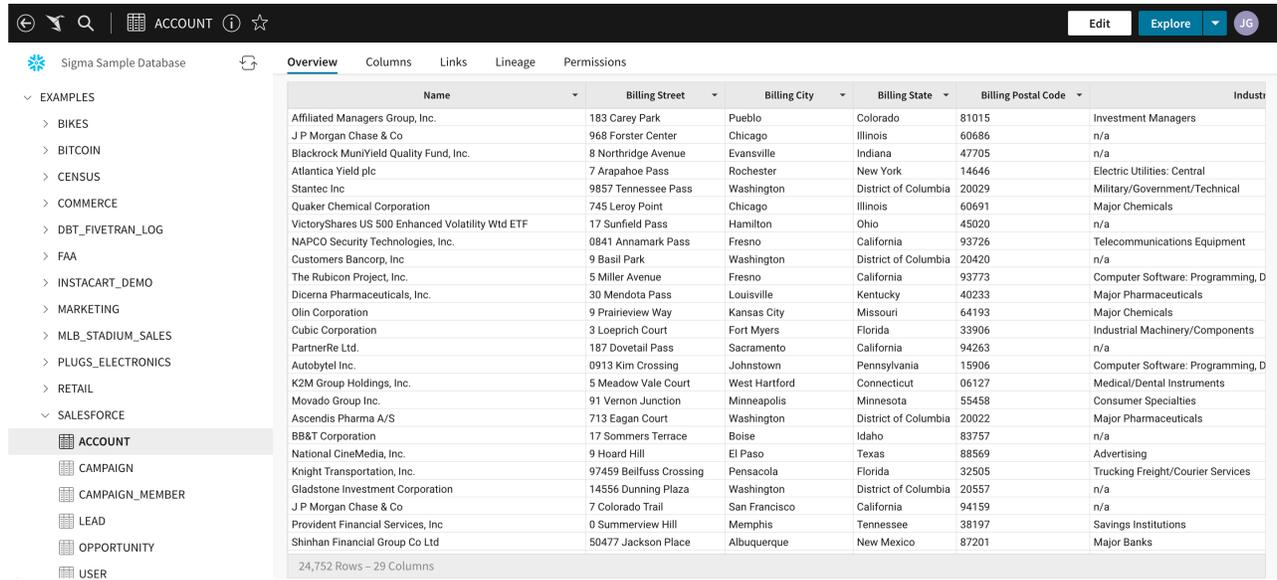


Figure 1: The interface of Sigma Workbooks connecting to a CDW. Consider a business user Joey who is interested in launching a sales campaign and identifying ideal customers. She first opens up the ACCOUNT table in their SALESFORCE database, which lists all customer companies. Next, she wants to augment the table with relevant columns to help make customer selections.

large number (million scale) of small tables. In contrast, we observe that Sigma customers’ CDWs have several thousand tables on average but can have billions of rows. Due to the relatively smaller number of tables but larger table sizes, index lookup time becomes less dominant compared to data loading time and profiling time, both of which can now bottleneck end-to-end query response time in the pipeline. So, the main challenge in the Sigma scenario shifts in part from searching over a massive number of small tables to optimizing the entire pipeline for a fair number of very large tables (e.g., reducing data loading time and using efficient profiling methods).

In this work, we introduce a data discovery system WARP GATE and share our progress. WARP GATE prototypes semantic join discovery in Sigma, an enterprise platform for cloud data analytics. Specifically, we make the following contributions:

- We present WARP GATE, a semantic join discovery system, which demonstrates a real-world usage scenario of data discovery within an enterprise product.
- We identify unique challenges in Sigma and propose an embedding-based solution to capture semantic relationships between tables, especially those across databases.
- We evaluate WARP GATE on three cross-domain table corpora, including one collected from a CDW, and show our embedding-based approach is effective and robust to sampling. We also discuss further optimization opportunities for deployment and larger-scale discovery.

2 PROBLEM DEFINITION

In this section, we first briefly introduce Sigma and describe the data discovery needs of Sigma users. Then we formally define our problem.

2.1 Data Discovery Need in Sigma

Sigma is a web-based SaaS platform for cloud data analytics and business intelligence. It features Sigma Workbooks, which provides a spreadsheet-like interface for business users to perform interactive visual analysis of data in CDWs.

Among many features familiar to business users who are more skilled in spreadsheet functions than SQL, Lookup enables users to write a spreadsheet formula and add a column from another data source (i.e., another table) to enrich the dataset at hand. Under the hood, Workbooks translates spreadsheet formulas into SQL queries required to fetch data from cloud databases and performs cardinality-preserving joins on related tables.

A major limitation of Lookup is that users have to specify the join path in formulas. In other words, it requires users to have prior knowledge about relationships between datasets. However, users often lack such knowledge and have to rely on their data teams for help or manually browse tables in databases to find a clue. This limitation has become a bottleneck in fully utilizing the Lookup functionality, significantly delaying their data exploration and analysis workflows. We believe this sort of data discovery will be a challenge for all cloud analytics products and hence has to be addressed.

2.2 Semantic Join Discovery

Problem Statement (Top- k Semantic Join Discovery). Given a corpus of tables \mathcal{S} , a query column c_q from a table Q , and a constant k , find up to k candidate columns from \mathcal{S} such that they are most likely to be joinable with c_q .

Traditionally, joinable tables have been discovered through syntactic analysis, which can miss many join opportunities, particularly

when data are independently sourced. High dimensional embeddings have recently been shown to capture semantics effectively in many domains [8, 23]. We expect that the semantic similarity of two columns can be quantified by some similarity measure between their embeddings, which we expect can serve as a proxy for their join-ability as well.

Definition (Semantic Column Join-ability). Given two columns A and B , we define semantic column join-ability $\mathcal{J}(A, B)$ as

$$\mathcal{J}(A, B) = \mathcal{M}(\mathcal{T}(A), \mathcal{T}(B))$$

where $\mathcal{T}(\cdot)$ is an embedding function that projects a given column into a vector space and $\mathcal{M}(\cdot, \cdot)$ is a similarity measure between two embeddings in the vector space.

Based on this metric, we can attempt to solve the semantic join discovery problem by returning candidate tables in descending order of semantic column join-ability \mathcal{J} relative to the query column c_q . How to do this effectively is the central challenge addressed in this paper.

3 WARP GATE OVERVIEW

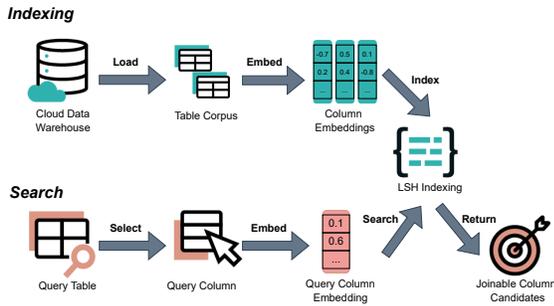


Figure 2: WARP GATE pipeline.

We discuss key challenges to implementing an embedding-based approach to the semantic join discovery problem, along with our solutions to these challenges in Section 3.1. Then we describe our system prototype WARP GATE incorporating these ideas in Section 3.2.

3.1 Embedding-Based Solution

Figure 2 gives an overview of the WARP GATE architecture, which consists of two main pipelines: indexing and search, which we consider in turn below. But first, we discuss the embedding step common to both pipelines.

3.1.1 Column Embeddings. We consider embeddings at the column level an effective encapsulation of semantics for join discovery. Intuitively, embeddings can capture semantic relationships between columns across databases where value formatting may vary widely despite semantic proximity.

However, the effectiveness of embeddings depends on the embedding model. We consider several factors in choosing an embedding model:

- (i) Whether the model is (pre-)trained over tabular data;
- (ii) The size of the training corpus;

- (iii) The efficiency of model inference.

We observe two primary paradigms in applying embeddings for data management tasks. Early work [4] leverages off-the-shelf embedding models from the natural language processing (NLP) field. Recent work [8, 17] adapts NLP model architectures and training algorithms for tabular data. As demonstrated in their work, embeddings from models trained on tabular data manifest better performance in downstream tasks. For instance, using the same training algorithm, [17] shows that models trained on sequences extracted from tables outperform the model trained on an unstructured text corpus. The performance uplift is attributed to the serialization of tabular data, which considers the table structure in the embedding process. With this observation, we prefer embedding models designed for tabular data.

Empirically, more training data usually lead to better model performance and generalizability, given enough model capacity. In the database community, Web table corpora [11, 18] extracted from the Common Crawl contribute the largest publicly accessible relational table datasets. Efforts exploring representation learning for relational tables [8, 17] have trained embeddings on various Web table corpus, which also exhibits the impact of Web tables beyond Web applications. Thus, we favor embedding models derived from those large table corpora.

The embedding process for tabular data generally involves table serialization, input encoding (e.g., tokenization and transformation), and model inference. It is worth noting that although more complex models (in terms of the number of parameters) give embeddings that tend to yield better results in the downstream tasks, they also carry a larger inference cost (i.e., longer inference time, which directly affects query response time) compared to simple models. For interactive applications such as join discovery and recommendation in Sigma, we need to choose models that balance the embedding quality and model inference time.

3.1.2 Indexing. We employ locality-sensitive hashing (LSH) [5, 16] to turn the high-dimensional embedding similarity search problem into nearest neighbor search in a low-dimensional space. The basic idea of LSH is to maximize hash collisions for similar inputs. With respect to a similarity measure, LSH applies a specific family of hash functions that assign two similar inputs to the same “buckets” with high probability that is equal to the similarity score between them.

In our case, we consider the semantic join-ability of two columns to be the cosine similarity of their embeddings. To approximate the cosine similarity of two vectors, we rely on SimHash (also known as random projection) [5] that essentially uses hyperplanes in the vector space as hash functions. We hash all column embeddings from the table corpus into a SimHash LSH index. When a query comes in, the LSH index hashes the query embedding and only searches within the sub-universe of embedding vectors (much smaller than the entire universe) that share the same hash as the query, hence reducing the search time.

3.1.3 Sampling. Existing data discovery systems [2, 3, 12] run a one-pass data scanning step to read all datasets once and retrieve profiling information (e.g., MinHash for each column). Reading full tables of millions of rows or even billions of rows from CDWs is

expensive in terms of both computation time and monetary cost because CDW vendors that adopt a usage-based pricing model (i.e., pay-as-you-go) charge per GB of data scanned. It is then intuitive to use sampling to reduce the cost. But one side effect of sampling is potential disruptions to column profiles. For instance, profiling methods like MinHash are shown to be sensitive to the sample size [6]. It also remains unclear what a good sample size is for embeddings to preserve the similarity for data discovery.

3.2 WarpGate Interface

We have implemented the ideas described above in a prototype system we call WARP_{GATE}. As shown in Figure 3, we integrate WARP_{GATE} into Sigma Workbooks as a functionality called Add column via lookup.

Continuing the example of a business user Joey who wants to identify customers for a sales campaign, she opens up the Account table and sees a full list of companies (the Name column highlighted in Step 1 of Figure 3). To help make the decision, she wants to augment the table with additional information so that she can know which companies fall into the scope of the sale campaign. In the current interface, she right-clicks on the Name column and chooses the Add column via lookup option in the menu. A window shows up (Step 2 in Figure 3) and displays top- k join path recommendations. Each recommendation gives the candidate join column, the table, and the database from which the candidate column comes, as well as a similarity score relative to the Name column (i.e., the query column). Although recommendations are ranked in descending order of the similarity score, Joey can pick one that fits her interest most. Once she chooses a candidate, the window will display a list of all the columns from the candidate table. Joey can browse and select the most relevant columns. For now, the interface will add user-selected columns to the Account table right at the side of the Name column for information complement and also load the candidate table into Sigma Workbooks for further inspection. In Step 3, we simply add and show the first recommended column, which is a foreign key in the same database.

The current interface design is based on the consideration that it is less desirable to simply join two wide tables with many columns and overwhelm users with a display of an even wider table that does not really help users with their information needs. In this case, Joey can first add the most relevant columns at first glance based on her domain knowledge while having the flexibility to see the entire candidate table for more careful examination.

4 EXPERIMENTS

We evaluate WARP_{GATE} on effectiveness and efficiency in comparison with two existing data discovery systems that support enterprise use cases. We obtain promising results of our embedding-based approach and show that it is sample-efficient and scales well to tables of millions of rows.

4.1 Datasets

We use three repositories for evaluation. Table 1 gives a summary of dataset characteristics.

Nextia_{JD}. [14] composes four testbeds of datasets from open repositories such as Kaggle and OpenML. Datasets are divided into

Table 1: Basic statistics of evaluation datasets. XS is shorthand for testbedXS in the Nextia_{JD} repository, similarly for S, M, and L.

	# Tables	# Columns	Avg. # Rows	# Queries	Avg. # Answers
XS	28	257	1,938	35	2.8
S	46	2,553	209,646	177	3.6
M	46	1,067	3,175,904	188	4.4
L	19	541	12,288,165	92	3.6
Spider	70	429	7632	60	1.1
Sigma	98	1,343	2,243,932	TBD	N/A

testbeds according to their file size. For instance, Nextia_{JD}-XS contains datasets of size smaller than 1 MB while Nextia_{JD}-L consists of datasets of size larger than 1 GB. They also label the join quality of pairs of attributes based on a measure that considers both containment and cardinality proportion with empirically determined thresholds. In experiments, we use attribute pairs with quality labeled as Good and High by [14].

Spider. Released as a large-scale semantic parsing and text-to-SQL dataset, Spider [22] includes 5,693 SQL queries on 200 databases across domains. We parse schema SQL files and retrieve join paths between primary keys and foreign keys as ground truth. Spider has both the training set and the development set, and we currently use join paths from the development set for evaluation.

Sigma. The Sigma Sample Database is a collection of schema and tables that are accessible to all Sigma Computing accounts for exploration. It is stored in a Snowflake account managed by Sigma Computing and changes over time. The datasets in this database vary in origin; some data is real, publicly accessible data, some is real-data that has been obfuscated, and some is completely auto-generated data. The domain of the data ranges includes retail (transactions, products and stores), financial (daily attributes of many securities), demographic (census, restaurants and bikes) and usage (cloud application usage and metering, server logs).

4.2 Baselines & Metrics

For comparison, we consider as baselines two system prototypes **Aurum** [12] and **D³L** [2] that report on real-world data discovery (see Section 6 for more information). Note that Aurum stores detected relationships in a graph data structure and does not support top- k search.

When ground truth is available, we report top- k precision and recall with varying k . At each value of k , we average the precision and recall numbers over all the queries. We also limit k to a small range since the average number of answers to a query is small in evaluation datasets, and we don't want to overwhelm users with too many recommendations.

Besides effectiveness, we consider index lookup time and end-to-end query response time for efficiency measurement. Index lookup time refers to the time for the index data structure to return top- k candidates, whereas end-to-end query response time is the total amount of time for a discovery system to respond to a query, which includes data loading and embedding inference time in our case. We report both metrics in seconds averaged over all queries in a dataset. It is crucial in our application scenario that the system responds to queries at interactive speed.

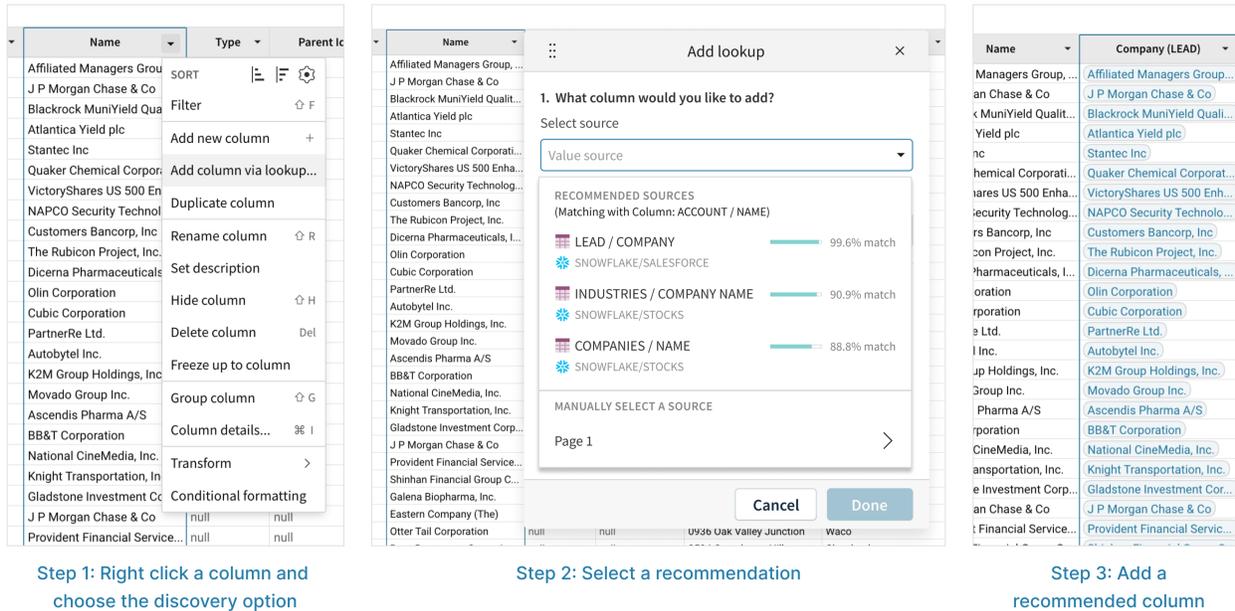


Figure 3: WARP_{GATE} is integrated into a functionality called Add column via lookup in Sigma Workbooks.

4.3 Experiment Results

Based on considerations in Section 3.1, we choose to use Web Table Embeddings [17] as our underlying embedding model. We set the similarity threshold of the SimHash LSH index to 0.7 and run baselines with their default setting. All experiments are conducted on an Amazon EC2 p3.8xlarge instance.

4.3.1 *Join Discovery on nextia_{JD} Testbeds.* We present discovery results on Nextia_{JD} testbedS and testbedM in Figure 4(a) and 4(b) respectively (results on four testbeds have similar trends while these two testbeds contain more tables and queries than the other two). As shown, WARP_{GATE} consistently obtains higher precision and recall compared with two baselines as *k* increases.

Table 2: End-to-end query response time in seconds/query (for *k* = 10) along with index lookup time reported in parentheses for WARP_{GATE}.

	Aurum	D ³ L	WARP _{GATE}
testbedS	0.18	4.77	3.12 (1.04)
testbedM	0.03	57.69	38.73 (8.39)

Table 2 shows end-to-end query response time. Aurum is significantly faster because it indexes detected relationships in another graph data structure for result retrieval. In contrast, D³L is the slowest since it uses an ensemble approach and aggregating multiple types of signals takes longer. As we can also see here, given the equal number of tables in two testbeds, when the average number of rows increases by an order of magnitude, query response time increases linearly and becomes unacceptable for interactive purposes. Moreover, index lookup time accounts for less than 25% of

query response time on testbedS and less than 13% on testbedM. This suggests that there are other bottlenecks in the pipeline and improving only index lookup time (as many earlier works focus on) is not enough to reduce overall query response time. Upon inspection, we found loading large datasets from disk to memory and embedding inference take most of the time besides index lookup. We show in Section 4.4 that sampling is an effective mitigation and our embedding approach is robust to sample size.

4.3.2 *PK/FK Detection on Spider.* Figure 4(c) shows the top-*k* precision and recall of three systems on the Spider dataset. We use this dataset mainly to demonstrate that for PK/FK detection within each independent database, the embedding measure alone (WARP_{GATE}) can compare favorably against the ensemble approach (D³L) and outperform the syntactic-only approach (Aurum) by a large margin. As to D³L, the jump of recall from *k* = 5 to *k* = 10 is attributed to the column name similarity measure in D³L, and we indeed observe in Spider that many PK/FKs share syntactically similar column names. Since Spider has a small number of tables that contain thousands of rows on average, the search time is fast for all three systems, which is no more than 2 seconds for running all the queries.

4.3.3 *Ad-Hoc Discovery in Sigma.* Since Sigma Sample Database is a corpus without ground truth, we ask four colleagues to pick columns of interest as queries. Here we describe one of the most interesting discoveries. In future work, we plan to conduct a user study to evaluate the value of WARP_{GATE} in a principled manner.

Continuing our running example of the business user Joey, who is interested in looking for additional information to help make the decision of customer selection. Figure 3 displays the top-3 recommendations from WARP_{GATE}. The first candidate is the Company column of the LEAD table from the same database as the query

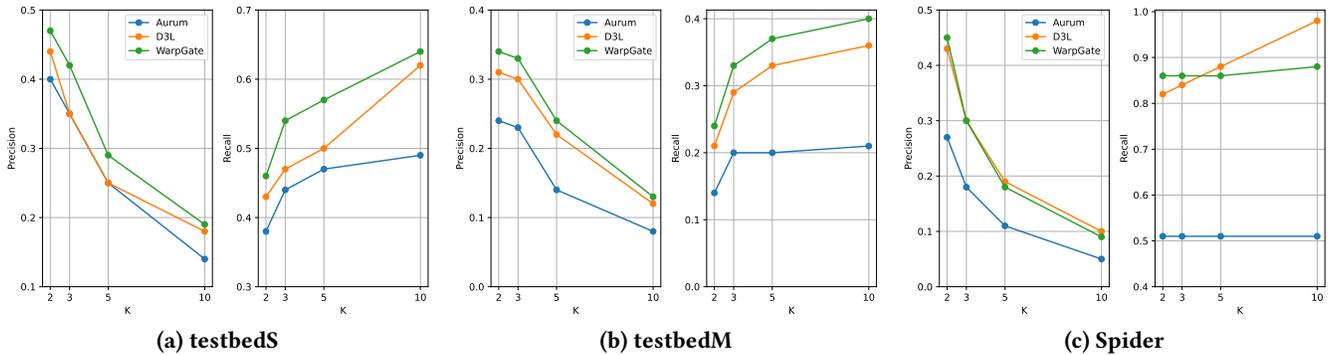


Figure 4: Top- k precision and recall.

column, which contains information on contact points in each company (e.g., name, title, and address of each contact point). Upon browsing columns in the LEAD table, Joey does not find auxiliary information for her need. She moves to inspect the second candidate, which is the Company Name column of the INDUSTRIES table from the STOCKS database. In the same table, the Industry Group column has sector information of US public companies, which can (partially) enrich the query table. Joey then acts on her customer selection by filtering companies with their sectors. Even more interestingly, she can add to the query table the TICKER column from the INDUSTRIES table, which can be further used as the join key to add columns of stock prices from the same STOCKS database. This way, Joey can track down high-performing companies in targeted sectors as her sales campaign customers.

4.4 Sample Efficiency

We run WARP GATE over Nextia_{JD}-S and Nextia_{JD}-M with sample size 10, 100, and 1000. For all sample sizes, the embedding approach remains as effective as using full values with $\pm 1\%/2\%$ variation at different k values. In the meanwhile, index lookup time is reduced by up to two orders of magnitude (e.g., from about 1 second to about 10 milliseconds on Nextia_{JD}-S), and end-to-end query response time also comes down to interactive speed (under 35 milliseconds per query on Nextia_{JD}-S and under 65 milliseconds per query on Nextia_{JD}-M). We discuss further optimization opportunities in Section 5.

To see if sample efficiency is unique to Web Table Embeddings, we also evaluate BERT [9] as our underlying embedding model. We observe that the effectiveness of BERT Embeddings is robust to sample size as well. Although BERT is a significantly more complex model, we do not see much improvement in effectiveness (mostly on par with Web Table Embeddings). Nevertheless, index lookup time and query response time are 10x slower without sampling due to more expensive embedding inference costs.

5 DISCUSSION

5.1 Sigma Customer Data Scale

The scale of enterprise CDWs is a challenge for any deployment of join discovery. The median Sigma customer warehouse has 450 tables for analysis, but the average is over 12,700 tables with an

average of 25.7 columns per table. The size of enterprise tables is also a challenge. The median table in these warehouses has just 7,700 rows, but the average is 1.7 billion rows. Actively sampling this large number of tables incurs usage costs in CDWs, so a passive sampling of user queries is beneficial. Likewise, samples can be shared across ML applications in Sigma to amortize their cost.

5.2 Optimization Opportunities

5.2.1 Contextual Embeddings. Existing approaches generally follow a two-step process: they first build profiles for each column independently and then determine join-ability of two columns based on some notion of similarity between their profiles (e.g., containment or Jaccard similarity). Typically, the context of a column is not taken into account when building those profiles. However, context (e.g., other columns in the same table, user activities, query logs) can potentially provide auxiliary information that is critical to find semantically related candidates. We plan to explore the option of incorporating context information into the underlying embedding model of WARP GATE.

5.2.2 Data Management for Data Discovery. CSV is a common file format for tabular data in Open Data Lakes and table repositories on the Web. Despite the flexibility and human readability of CSV files, the flat-file format is not ideal for storage and query purposes. We note that (1) loading giant CSV files and running embedding inference can incur heavy memory usage and crash programs on small machines; (2) operations in join discovery like embedding inference are column-oriented; (3) what is unique in our scenario is that our data are pulled from CDWs, thus more structured and much better data quality than those in the wild. (1) and (2) make us think it is more memory and storage efficient to take advantage of (in-memory) column stores for data pulled out of CDWs, and (3) can make the process easier. This change will also have significant implications for deployment as we need to provision CPU and memory usage in Sigma’s Kubernetes cluster.

5.2.3 Efficient Search. To further improve query response time for Sigma customers with a massive number of tables in their CDWs, we are considering two directions: (1) a block-and-verify strategy in [10], which employs pivot-based filtering to prune dissimilar vectors; (2) fine-tune off-the-shelf embedding models in a self-supervised way that pushes embeddings of joinable columns to

have higher cosine similarity so that an index data structure like SimHash can be better utilized [7].

6 RELATED WORK

Data Discovery Systems for Enterprise. Aurum [12] presents a data discovery system that models syntactic relationships between datasets in a graph data structure. In a two-step process of profiling data and computing the relationship between data signatures, Aurum builds up a graph with nodes representing column profiles and weighted edges indicating syntactic relationships between two nodes (e.g., content similarity). With an in-memory index of the graph, Aurum can efficiently support varied discovery needs such as how similar two columns are, which can be used to find joinable datasets. SEMPROP [13] further employs word embeddings in Aurum to surface semantically related objects in the graph.

Voyager [3] built on top of D³L formalizes and supports three tasks routinely performed by data scientists using their platform. It profiles data with a full scan and indexes columns with MinHash. Voyager also drops the embedding measure in D³L and relies only on syntactic measures.

Join Discovery over Open Data. JOSIE [25] considers the problem of finding joinable tables in modern data lakes. They formulate the problem as a top-k overlap set similarity search problem where columns are treated as sets and matching values as intersection between sets. By using a cost model for estimating the reading costs of data structures, JOSIE contributes an exact overlap set similarity search algorithm that can adapt to different data distributions.

JUNEAU [24] targets on search and management problems including finding joinable tables in interactive data science environments (e.g., Jupyter Notebook). It combines multiple table-relatedness measures (e.g., column and row overlap, provenance and etc.) for related table search.

D³L [2] proposes a distance-based framework that uses five types of evidence to determine column similarity: (i) column name similarity; (ii) column extent overlap; (iii) word-embedding similarity; (iv) format representation similarity; and (v) domain distribution similarity for numerical attributes. Their aggregated approach is shown to be more effective and efficient than Aurum on table union search, which is a problem that has significant overlapping with join path discovery.

In contrast to the search paradigm, [20] proposes building a navigation graph over data lakes. By leveraging a probabilistic model, they show that navigation can help users find relevant tables that cannot be found by keyword search.

7 CONCLUSION

We present WARP GATE, a system prototype for semantic join discovery in Sigma Workbooks. In future work, we expect to share our experience and lessons of deploying WARP GATE in the production environment and conduct a user study to further evaluate its value for end users.

ACKNOWLEDGMENTS

We thank Siva Obulam, Jade Fleishhacker, and Qing Feng for contributing to the first demonstration of WARP GATE in Sigma Summer Hackathon 2022. We thank Oscar Bashaw for sharing with us the

Sigma Sample Database and his description of this corpus. We also thank Nipurn Doshi for designing the WARP GATE interface. This project is supported in part by NSF grant 1946932.

REFERENCES

- [1] Sagar Bharadwaj, Praveen Gupta, Ranjita Bhagwan, and Saikat Guha. 2021. Discovering Related Data At Scale. *Proc. VLDB Endow.* (2021), 1392–1400.
- [2] Alex Bogatu, Alvaro A. A. Fernandes, Norman W. Paton, and Nikolaos Konstantinou. 2020. Dataset Discovery in Data Lakes. In *ICDE 2020*. IEEE, 709–720.
- [3] Alex Bogatu, Norman W. Paton, Mark Douthwaite, and André Freitas. 2022. Voyager: Data Discovery and Integration for Onboarding in Data Science. In *EDBT 2022*. OpenProceedings.org, 2:537–2:548.
- [4] Rajesh Bordawekar and Oded Shmueli. 2017. Using Word Embedding to Enable Semantic Queries in Relational Databases. In *DEEM@SIGMOD 2017*. ACM, 5:1–5:4.
- [5] Moses Charikar. 2014. Similarity estimation techniques from rounding algorithms. In *STOC 2014*. ACM, 380–388.
- [6] Zhimin Chen, Vivek R. Narasayya, and Surajit Chaudhuri. 2014. Fast Foreign-Key Detection in Microsoft SQL Server PowerPivot for Excel. *Proc. VLDB Endow.* (2014), 1417–1428.
- [7] Tianji Cong, Fatemeh Nargesian, and H. V. Jagadish. 2022. Pylon: Semantic Table Union Search in Data Lakes. *Under review*.
- [8] Xiang Deng, Huan Sun, Alyssa Lees, You Wu, and Cong Yu. 2020. TURL: Table Understanding through Representation Learning. *Proc. VLDB Endow.* (2020), 307–319.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT 2019*. Association for Computational Linguistics, 4171–4186.
- [10] Yuyang Dong, Kunihiro Takeoka, Chuan Xiao, and Masafumi Oyamada. 2021. Efficient Joinable Table Discovery in Data Lakes: A High-Dimensional Similarity-Based Approach. In *ICDE 2021*. IEEE, 456–467.
- [11] Julian Eberius, Maik Thiele, Katrin Braunschweig, and Wolfgang Lehner. 2015. Top-k entity augmentation using consistent set covering. In *SSDBM 2015*. ACM, 8:1–8:12.
- [12] Raul Castro Fernandez, Ziawasch Abedjan, Famiem Koko, Gina Yuan, Samuel Madden, and Michael Stonebraker. 2018. Aurum: A Data Discovery System. In *ICDE 2018*. IEEE Computer Society, 1001–1012.
- [13] Raul Castro Fernandez, Essam Mansour, Abdulhakim Ali Qahtan, Ahmed K. Elmagarmid, Ihab F. Ilyas, Samuel Madden, Mourad Ouzzani, Michael Stonebraker, and Nan Tang. 2018. Seeping Semantics: Linking Datasets Using Word Embeddings for Data Discovery. In *ICDE 2018*. IEEE Computer Society, 989–1000.
- [14] Javier Flores, Sergi Nadal, and Oscar Romero. 2021. Towards Scalable Data Discovery. In *EDBT 2021*. OpenProceedings.org, 433–438.
- [15] James Gale, Max Seiden, Deepanshu Utkarsh, Jason Frantz, Rob Woollen, and Çagatay Demiralp. 2022. Sigma Workbook: A Spreadsheet for Cloud Data Warehouses. *Proc. VLDB Endow.* (2022), 3670–3673.
- [16] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *VLDB 1999*. 518–529.
- [17] Michael Günther, Maik Thiele, Julius Gonsior, and Wolfgang Lehner. 2021. Pre-Trained Web Table Embeddings for Table Discovery. In *aiDM@SIGMOD 2021*. ACM, 24–31.
- [18] Oliver Lehmborg, Dominique Ritze, Robert Meusel, and Christian Bizer. 2016. A Large Public Corpus of Web Tables containing Time and Context Metadata. In *WWW 2016*. ACM, 75–76.
- [19] Renée J. Miller, Fatemeh Nargesian, Erkang Zhu, Christina Christodoulakis, Ken Q. Pu, and Periklis Andritsos. 2018. Making Open Data Transparent: Data Discovery on Open Data. *IEEE Data Eng. Bull.* (2018), 59–70.
- [20] Fatemeh Nargesian, Ken Q. Pu, Erkang Zhu, Bahar Ghadiri Bashardoost, and Renée J. Miller. 2020. Organizing Data Lakes for Navigation. In *SIGMOD 2020*. ACM, 1939–1950.
- [21] Fatemeh Nargesian, Erkang Zhu, Ken Q. Pu, and Renée J. Miller. 2018. Table Union Search on Open Data. *Proc. VLDB Endow.* (2018), 813–825.
- [22] Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir R. Radev. 2018. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *EMNLP 2018*. Association for Computational Linguistics, 3911–3921.
- [23] Dan Zhang, Yoshihiko Suhara, Jinfeng Li, Madelon Hulsebos, Çagatay Demiralp, and Wang-Chiew Tan. 2020. Sato: Contextual Semantic Type Detection in Tables. *Proc. VLDB Endow.* (2020), 1835–1848.
- [24] Yi Zhang and Zachary G. Ives. 2020. Finding Related Tables in Data Lakes for Interactive Data Science. In *SIGMOD 2020*. ACM, 1951–1966.
- [25] Erkang Zhu, Dong Deng, Fatemeh Nargesian, and Renée J. Miller. 2019. JOSIE: Overlap Set Similarity Search for Finding Joinable Tables in Data Lakes. In *SIGMOD 2019*. ACM, 847–864.