



# Unveiling Challenges for LLMs in Enterprise Data Engineering

Jan-Micha Bodensohn\*  
DFKI & Technical University of  
Darmstadt

Ulf Brackmann\*  
SAP SE & DFKI

Liane Vogel\*  
Technical University of Darmstadt

Anupam Sanghi†  
Indian Institute of Technology  
Hyderabad

Carsten Binnig  
Technical University of Darmstadt  
& DFKI

## ABSTRACT

Large Language Models (LLMs) promise to automate data engineering on tabular data, offering enterprises a valuable opportunity to cut the high costs of manual data handling. But the enterprise domain comes with unique challenges that existing LLM-based approaches for data engineering often overlook, such as large table sizes, more complex tasks, and the need for internal knowledge. To bridge these gaps, we identify key enterprise-specific challenges related to data, tasks, and background knowledge and extensively evaluate how they affect data engineering with LLMs. Our analysis reveals that LLMs face substantial limitations in real-world enterprise scenarios, with accuracy declining sharply. Our findings contribute to a systematic understanding of LLMs for enterprise data engineering to support their adoption in industry.

## PVLDB Reference Format:

Jan-Micha Bodensohn, Ulf Brackmann, Liane Vogel, Anupam Sanghi, and Carsten Binnig. Unveiling Challenges for LLMs in Enterprise Data Engineering. PVLDB, 19(2): 196 - 209, 2025.

doi:10.14778/3773749.3773758

## PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/DataManagementLab/llmefl-enterprise-challenges>.

## 1 INTRODUCTION

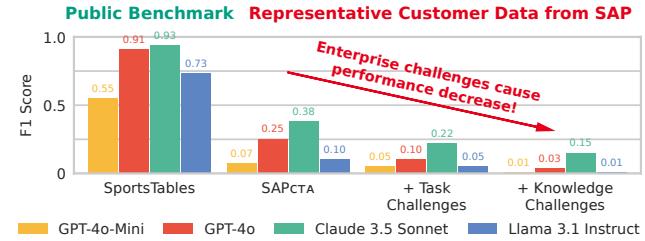
Large enterprises generate vast amounts of tabular data that drives applications like machine learning and analytical query processing. Data engineering is crucial for understanding this raw data and preparing it for its downstream usage. It encompasses a range of tasks, from data exploration and integration to data transformation and cleaning. Since these tasks often impose significant manual overhead to apply existing tools to the specific data at hand, the automation of individual data engineering tasks like entity matching [28, 40] and column type annotation [21, 69] with the help of machine learning has long drawn attention from researchers. Nevertheless, adapting such machine learning approaches to new

\*Authors with equal contribution, alphabetical ordering. Correspondence goes to jan-micha.bodensohn@cs.tu-darmstadt.de

†Work done while the author was at Technical University of Darmstadt.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 2 ISSN 2150-8097.  
doi:10.14778/3773749.3773758



**Figure 1: LLMs perform well on public benchmarks but poorly in real-world enterprise settings.** The plot compares support-weighted F1 scores for column type annotation on the public SportsTables benchmark with customer data from SAP. We further raise the difficulty (*+ Task Challenges*) by increasing the number of semantic types from 200 (comparable to public benchmarks) to the full 5,089 from the enterprise setting and observe an additional performance drop. When also requiring internal knowledge about company-specific schema extensions in the form of customer-defined columns (*+ Knowledge Challenges*), the performance is close to zero.

datasets and tasks often requires computer science expertise, rendering them inaccessible to many practitioners.

Recent work has shown that Large Language Models (LLMs) can be directly applied to data engineering tasks on tabular data, indicating that they achieve state-of-the-art results on various table-based tasks without requiring task-specific architectures and training [3, 22, 25, 41]. Their out-of-the-box nature provides a significant advantage over other machine learning approaches that require supervised training for each dataset and task. One example is the task of column type annotation, where the goal is to annotate the columns of a relational table with semantic types from a given ontology. Whereas machine learning-based approaches like Sherlock [21] and Sato [69] require re-training for each new set of semantic types, LLMs can easily support different sets of semantic types by including them in the prompt [29]. Therefore, they provide a promising avenue to automate data engineering tasks.

**Enterprise challenges.** Some first papers have now made the point that while LLMs achieve strong results on existing benchmarks, they fall short when applied to real-world enterprise data [6, 26]. A central observation here is that the data in public benchmarks often comes from web resources like Wikipedia [4] and GitHub [20]. In contrast, the data from companies running their business processes with software systems like those from SAP differs fundamentally

from these datasets in many aspects, including table sizes, sparsity, and data types [26, 54, 61]. Since LLMs are typically trained on public data scraped from the web [7, 39], they have not seen much of this enterprise data during their training.

In this paper, we set out to systematically study the performance of LLMs for enterprise data engineering on representative customer data and case studies reflecting real enterprise scenarios. As a sneak peek into our results, Figure 1 presents a first experiment comparing the performance of LLMs for column type annotation on a public benchmark to a dataset of real-world customer data from SAP. As shown by the left bar group, LLMs of different types and sizes achieve high F1 scores when evaluated on a public benchmark like the SportsTables dataset [31]. The same models, however, show substantial performance decreases when applied to the representative customer data from SAP for column type annotation with a comparable number of semantic types, as shown in Figure 1 (*SAPCTA*).

Aside from these data challenges, data engineering in enterprises comes with additional difficulties:

(1) *Task complexity*: Enterprise tasks are often more complex than their academic formulations [54]. One example is the task of entity matching, where existing literature often assumes that each entity is one row in a single table [28, 49]. In enterprises, however, entities are often business objects that span across multiple tables, making matching difficult. Moreover, the tasks themselves are often more complex. For example, increasing the complexity of column type annotation by scaling to the true number of semantic types (5,089 instead of 200) in the SAP system leads to a further decrease in F1 scores, as shown in Figure 1 (+ *Task Challenges*).

(2) *Internal knowledge*: Data engineering in enterprises also often requires internal knowledge that is absent from public sources, limiting LLMs’ abilities to understand the data without additional context. This is especially true for schema customizations that involve customer-defined semantic types. On SAP’s column type annotation dataset, when adding those columns to the task, F1 scores for such customer-defined columns are near zero, as shown in Figure 1 (+ *Knowledge Challenges*).

We want to highlight how such enterprise-specific challenges affect LLMs. Our findings go beyond observations already known to the community and help unmask the jagged out-of-the-box capabilities of LLMs. For example, the impact of the complex enterprise schemas and symbolic values in the enterprise data is particularly striking. Moreover, we were surprised by how strongly the models resisted overwriting their parametric knowledge despite explicit guidance from documentation and examples. We believe this to be the first attempt to examine LLMs for data engineering on real enterprise data at this breadth. We see it as an important first step (of many) to make LLMs viable for enterprise data engineering.

**Contributions.** Our main contributions are:<sup>1</sup> (1) We systematically analyze the challenges involved in enterprise data engineering and structure them along the dimensions of enterprise *data*, enterprise *tasks*, and enterprise *knowledge*. (2) We experiment on representative enterprise data to show how it differs from existing public benchmarks and understand how it affects LLMs. (3) We

**Table 1: Model characteristics.** We evaluate five LLMs from three model providers covering multiple types and sizes.

	Context Window	Reason-ing	USD Per 1M	
			Input	Output
<b>GPT-4o-Mini</b> (2024-07-18) <sup>1</sup>	128K	no	0.15	0.60
<b>GPT-4o</b> (2024-08-06) <sup>1</sup>	128K	no	2.50	10.00
<b>o1</b> (2024-12-17) <sup>1</sup>	200K	yes	15.00	60.00
<b>Claude 3.5 Sonnet</b> (v1/v2) <sup>2</sup>	200K	no	3.00	15.00
<b>Llama 3.1 Instruct</b> (70B) <sup>3</sup>	128K	no	0.72	0.72

Pricing by OpenAI<sup>1</sup>, Anthropic<sup>2</sup>, and AWS Bedrock<sup>3</sup> in February 2025.

conduct multiple case studies that reflect real-world enterprise scenarios, allowing us to evaluate how different challenges affect LLMs in isolation. (4) We discuss directions for addressing these challenges, as well as the costs of using LLMs at enterprise scale. (5) Our experiments are performed with five recent LLMs from OpenAI, Anthropic, and Meta. To enable follow-up research, we make the code—including the full evaluation setup, all prompts, and where legally possible also the data—of this paper available to the broader research community.

## 2 SETTING OF OUR STUDY

Using LLMs to solve table-based tasks is a promising research direction that has been actively studied in recent years [29, 41, 49, 70, 71]. This section introduces the models we evaluate and briefly summarizes existing research on data engineering with LLMs.

**Choosing LLMs for the study.** LLMs are text foundation models trained on large corpora to complete natural language inputs [7] and follow user instructions [36, 63]. Recent LLMs like GPT-4 [45] and Llama 3 [16] build on a range of further ideas, from mixture-of-experts architectures [55] to more efficient implementations [9], leading to improvements in abilities like applying background knowledge and handling long inputs. With the release of reasoning LLMs like OpenAI’s o1 [43] and Deepseek’s R1 [10], they perform even better on tasks requiring several intermediate steps.

While many LLMs would be relevant candidates to include in our evaluation, our selection is necessarily limited. To ensure generalizable results, we evaluate five recent LLMs from three model providers covering multiple types and sizes and including open and closed models. As shown in Table 1, we use GPT-4o [42] as the state-of-the-art and GPT-4o-Mini as the cheapest model from OpenAI in February 2025, and o1 [44] as a reasoning model. We further include Claude 3.5 Sonnet [1] from Anthropic and Llama 3.1 Instruct [39] from Meta. In general, LLMs bear the challenge of output variability. In all our experiments, we reduce randomness to a minimum by setting the sampling temperature to 0. We also repeat several of our experiments multiple times, witnessing only small standard errors and no change in the overall characteristics of the results.

**Positioning the study.** Although several papers have already observed the differences between enterprise and web data [24, 54, 61, 72], existing research on data engineering with LLMs primarily uses evaluation datasets based on tables from public web resources, calling the applicability of LLMs on real-world enterprise

<sup>1</sup>This paper extends our previously published work [5, 6] on this topic.

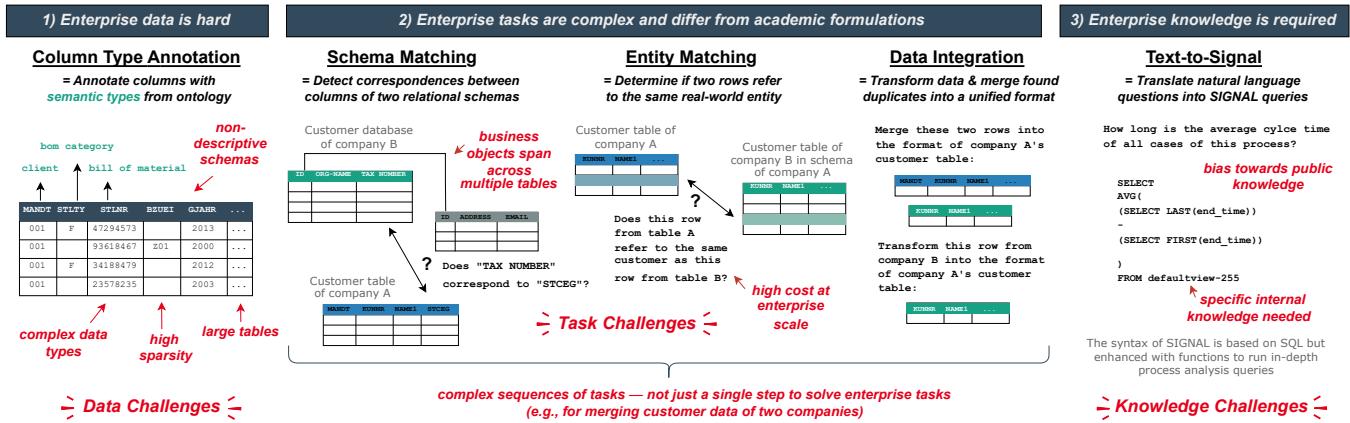


Figure 2: Enterprise-specific challenges in data engineering tasks. We use five well-established tasks serving as examples to highlight the breadth of challenges in enterprise settings and show their effect on LLMs for data engineering. The challenges shown here (e.g., high data sparsity) are highly general and extend to many other tasks beyond the ones shown here.

data into question. More recently, some enterprise-specific benchmarks have been released for tabular prediction tasks [27] and Text-to-SQL [8, 47]. Related to our work, Kayali et al. [26] quantify performance gaps between private and public data for the task of column type annotation and find that benchmarks based on public data overestimate the performance of LLMs. However, these studies focus only on individual tasks and consider the data as the only challenge in enterprise scenarios. In contrast, we perform case studies on various tasks to systematically analyze the challenges along several dimensions of working in enterprise scenarios.

**Working with real enterprise data.** Attempts like ours often fail because enterprise data is usually highly confidential and, therefore, hard to use in evaluations. For this paper, we were able to experiment with actual customer data from the enterprise systems of SAP. While one could argue that this study is still limited because we only use data from SAP, and there are clearly many alternative enterprise systems, SAP stands out as a dominant player in enterprise software systems across multiple industries worldwide. As such, we believe that our insights based on SAP data are highly valuable on their own and hope that our paper inspires other researchers with access to similar enterprise datasets to repeat our evaluations on their data. Importantly, the core characteristics of SAP data reflect the findings from other papers observing the differences between enterprise and web data [24, 54, 61, 72], highlighting the generality of our results beyond SAP data.

### 3 DESIGN OF OUR STUDY

We aim to take a holistic view of data engineering in enterprises by systematically analyzing the out-of-the-box performance of LLMs, covering enterprise-specific challenges along the dimensions of *data*, *tasks*, and *knowledge*. Figure 2 provides an overview of the tasks we have chosen to analyze the performance along each dimension. Below, we explain the rationale for selecting these tasks as examples to examine challenges that arise broadly across different enterprise tasks and scenarios. Sections 4 to 6 then describe our experimental evaluation for each of these challenges.

**The data challenge.** Enterprise data differs from public evaluation datasets in various aspects. First, the tables in enterprise databases are substantially larger in their number of columns and rows, and the schema and data are often more complex with column names and values that lack intuitive meaning. Moreover, enterprise data exhibits a much higher sparsity than public data, leaving many cells empty. Together, these factors make data engineering on enterprise data much more challenging. To explore these challenges, we focus on the task of column type annotation (Figure 2 left) [21, 69]. We have chosen this task because it is well-studied in the literature and thus enables the comparison of accuracies with public datasets. Moreover, the task can be formulated in various ways, stressing different data factors (e.g., with and without schema information) and allowing us to study the effects of different challenges like table size, sparsity, and descriptiveness of the data and schema.

**The task challenge.** Beyond the complexities of enterprise data, the enterprise tasks themselves are also more complex. First, while academia typically studies data engineering tasks in isolation, enterprise tasks in practice are often compounds of multiple simpler data engineering steps. As such, we argue that the evaluation procedures themselves must change and we need to study accuracy end-to-end to analyze effects such as how errors propagate. In this paper, we examine the compound nature of enterprise tasks using the example of integrating two customer databases, as illustrated in Figure 2 (middle). This process involves the steps of schema matching, entity matching, and data integration. A second challenge for enterprise tasks (not shown in Figure 2) is that even the individual steps are often more complex. For example, while entity matching in academia assumes 1:1 matches across rows of two tables [40, 50], matching in enterprise scenarios covers more complex cases like mapping several bank transfers to one invoice, and the required data is often scattered across multiple tables.

**The knowledge challenge.** As a last challenge, data engineering in enterprises often requires enterprise-specific knowledge. This is particularly challenging because LLMs will likely not have seen the required information during training, as it is covered only in

**Table 2: Data characteristics of publicly available benchmark datasets compared to representative customer data from SAP. Enterprise tables have substantially more rows and columns and display a higher sparsity compared to public data. Although most attributes are of type NVARCHAR, the data is highly symbolic, and table names, column names, and cell values are often not human readable because of abbreviations and enterprise-specific encodings.**

Tables	Columns		Rows		Sparsity <sup>1</sup>	Data Types <sup>2</sup>		Column Type Annotation	
	Med	95th	Med	95th		abc	123	Column Types	Labeled Columns
WikiTables-TURL	397,098	1	3	8	43	0.12	1.00	0.00	255
SOTAB	59,548	7	17	33	721	0.08	0.85	0.15	91
GitTablesCTA	1,100	12	33	25	263	0.12	0.33	0.67	122   59 <sup>3</sup>
SportsTables	1,183	21	31	32	924	0.07	0.16	0.84	452
<b>SAPCTA</b>	100 <sup>4</sup>	46	343	473,038	50,836,964	0.43	0.55	0.45 <sup>5</sup>	5,089
									8,106

<sup>1</sup> Sparsity is the fraction of empty cells. <sup>2</sup> Non-numerical (abc) and numerical (123) columns determined by pandas. <sup>3</sup> Using semantic types from DBpedia | Schema.org.

<sup>4</sup> We experiment on a representative sample from the thousands of tables in the customer system. <sup>5</sup> Only 14% of columns have numerical SQL types like INT and DECIMAL.

internal documentation or even just certain implementation details in the code of enterprise systems. To analyze these challenges, we have selected a task related to data exploration that requires translating natural language queries about business processes into the enterprise-specific query language SIGNAL, which differs slightly from SQL (Figure 2 right) [23]. The task is interesting because LLMs should, judging by its similarity to Text-to-SQL, in principle be able to solve it. However, they lack particular information about how exactly SIGNAL differs from SQL. At the core, the question we want to answer here is to what extent the lack of enterprise knowledge affects the accuracy of LLMs, and whether it can be increased by providing enterprise-specific knowledge as context.

**An orthogonal challenge: cost.** Since enterprise tables can contain millions of rows, using LLMs on such data can cause high costs, rendering some of the larger, more complex LLMs economically unviable. We discuss this aspect in Section 7.

## 4 THE DATA CHALLENGE

In this section, we quantify the anatomy of enterprise data by comparing real-world customer data from SAP to publicly available table corpora. We point out four challenges specific to enterprise data and perform experiments to evaluate how they affect LLMs, using column type annotation as an example task (Figure 2 left).

### 4.1 Data Challenges

For our study, we constructed a new corpus called SAPCTA. Table 2 compares the data characteristics of the real-world customer data in SAPCTA to several publicly available column type annotation benchmarks: WikiTables-TURL [11], SOTAB [30], GitTablesCTA [19], and SportsTables [31]. We observe the following differences:

**C1: Table size.** A first important observation is that enterprise tables often have substantially more rows and columns than the tables in public corpora. As shown in Table 2, some tables have hundreds of columns and millions of rows. While the large scale is a well-established data management problem [72], it poses challenges for LLMs, which have limited context windows. Although recent models have extended context windows, feeding large tables into LLMs still has downsides since latency and cost depend on the input size, and recent studies have shown that long contexts can lead to degraded performance for data residing in the “middle” [35].

**C2: Descriptiveness.** Another important insight is that schema properties like table and column names are often not descriptive but rather abbreviations that can only be understood with background knowledge or additional metadata [24]. This additional metadata is often unavailable or may not fit into the context window of the LLM. Moreover, the background knowledge is often specific to the particular enterprise, causing challenges for LLMs trained exclusively on publicly available data, as we discuss in Section 6.

**C3: Sparsity.** A third insight is that enterprise data is highly sparse. Table 2 shows that on average, 43% of the cells in enterprise tables are empty, compared to only 7-12% in existing datasets. The high sparsity results in a significant lack of information, posing a challenge for LLMs which rely on contextual cues to make accurate predictions. Moreover, we find that in addition to empty values, the cells in enterprise tables often contain dummy values such as 00000 that also denote the absence of an actual value.

**C4: Data types.** Interestingly, we find that only 45% of the columns in SAPCTA are classified as numerical by pandas, suggesting that enterprise datasets can be more text-heavy than found in previous studies [31]. Moreover, only 14% of the columns in the database schema have numerical data types like INT and DECIMAL. A closer inspection of the actual data reveals that the non-numerical data type NVARCHAR is often used to store symbolic values and codes such as invoice and material numbers, which is in line with previous findings [61]. Since these values are not self-expressive, LLMs cannot make use of them without additional context.

### 4.2 Experiments & Results

To study how the challenges C1-C4 affect LLMs for data engineering, we compare the models’ performance on our real-world enterprise dataset SAPCTA to their performance on existing evaluation datasets.

*Task.* We evaluate on column type annotation, a well-established data engineering task where the goal is to annotate the columns of a relational table with semantic types from a pre-defined ontology [21, 69]. We see it as an interesting example task to uncover the challenges of understanding enterprise tables with LLMs, since it requires a semantic understanding of the content of each column as well as the values of other columns and the table schema, all of which provide important signals to derive a semantic type.

**Table 3: Enterprise vs. public tables.** The table shows support-weighted F1 scores for column type annotation with and without column names. The results on enterprise data are substantially worse than on existing benchmarks.

Column Names	GitTablesCTA		SportsTables		SAPCTA	
	w/out	with	w/out	with	w/out	with
GPT-4o-Mini	0.52	0.96	0.27	0.55	0.02	0.07
GPT-4o	0.56	<b>0.99</b>	0.57	0.91	0.04	0.24
Claude 3.5 Sonnet	<b>0.67</b>	0.96	<b>0.66</b>	<b>0.93</b>	<b>0.05</b>	<b>0.34</b>
Llama 3.1 Instruct	0.49	0.95	0.38	0.73	0.02	0.10

**Setup.** Our SAPCTA corpus spans diverse business domains such as Finance, Sales and Distribution, Material Management, and Production Planning. For our experiments, we select 100 representative tables from the larger corpus, which contains multiple thousands of tables. The selection is based on discussions with SAP experts to identify the most widely used tables in the system. In total, the dataset includes 5,089 semantic types, such as *Amount Difference in Local Currency* or *Product Cost Collector*, which are generally more fine-granular than those used in other datasets [19, 30].

Our prompting strategy builds on best practices from existing literature, where the model annotates the columns of a given table based on a list of semantic types and one randomly-selected example included in the prompt [29, 69]<sup>2</sup>. Based on ablations (see Appendix A.1 in our extended technical report), we limit each table to three randomly-selected rows as a reasonable trade-off between performance and cost. For similar reasons, we serialize the tables in CSV format, which requires fewer formatting tokens than other serialization schemes like Markdown and JSON [56, 57]. We instruct the model to generate the column types as a JSON-formatted list.

**Exp. 1: Enterprise vs. public tables.** In our first experiment, we compare the performance for column type annotation on our SAPCTA corpus with the performance on GitTablesCTA [19] and SportsTables [31]. We perform each experiment twice, with and without including the table and column names (i.e., the table schema) in the prompt. Existing evaluations typically leave out the column names, since the semantic types are directly derived from them and the task would thus become trivial. For SAPCTA, however, the task is much harder. Therefore, we want to investigate how much the additional information helps.

Table 3 shows the results of this experiment. We make the following observations: First, LLMs have severe issues with column type annotation on enterprise data, leading to substantially worse results compared to the public benchmarks GitTablesCTA and SportsTables. The results are particularly poor in the experiments without table and column names, indicating that enterprise data on its own contains few helpful signals. Adding table and column names to the prompt improves the results, but they still remain much lower than for web tables. The remaining performance gap could potentially be attributed to the non-descriptive schema, the extremely wide and sparse tables, and the complex data types.

<sup>2</sup>To focus solely on the data challenges, we include only a subset of the semantic types in every prompt similar to the setup from Figure 1 (SAPCTA).

**Table 4: Non-numerical (abc) vs. numerical (123) data.** The table shows support-weighted F1 scores for column type annotation with column names. The results on numerical enterprise data are consistently worse.

Data Types	GitTablesCTA		SportsTables		SAPCTA	
	abc	123	abc	123	abc	123
GPT-4o-Mini	0.97	0.95	0.68	0.53	0.11	0.03
GPT-4o	<b>0.99</b>	<b>0.98</b>	<b>0.87</b>	0.91	0.31	0.16
Claude 3.5 Sonnet	0.97	0.95	0.81	<b>0.96</b>	<b>0.41</b>	<b>0.27</b>
Llama 3.1 Instruct	0.94	0.96	0.85	0.72	0.15	0.05

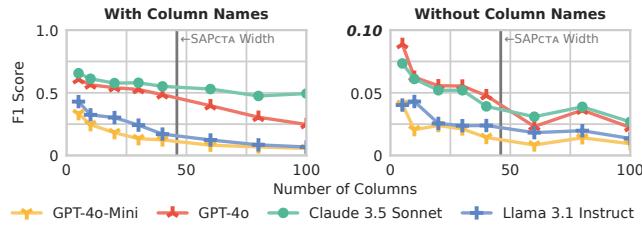
**Exp. 2: Textual vs. numerical data.** LLMs are known to often perform better on textual data than on numerical data [13, 31]. To study this effect on enterprise data, we compare the performance for non-numerical and numerical columns in our SAPCTA corpus.

Table 4 shows that, as expected, we see a higher performance on non-numerical enterprise data. By contrast, the results on the public benchmarks GitTablesCTA and SportsTables remain inconclusive, with only small performance gaps between non-numerical and numerical columns that also differ between models. The low performance on the SAPCTA dataset indicates that numerical data in enterprise systems is even harder to understand than in public benchmarks. Furthermore, the low scores for non-numerical columns may stem from the fact that enterprise tables often store identifiers like INV0014056 as type NVARCHAR.

**Exp. 3: Table width and sparsity.** To further investigate the performance gap between public benchmarks and enterprise data, we incrementally adapt the enterprise tables to resemble the characteristics of web tables more closely. Since two of the main differences are table width and sparsity, we vary the number of columns per table by randomly sampling subsets of columns and vary the sparsity by initially selecting only non-sparse columns and randomly removing individual cell values.

Figure 3 shows that increasing table widths lead to substantially worse results, indicating that the large table widths in enterprise data are indeed a major problem for LLMs. Figure 4 shows that increased sparsity leads to worse results only if no table and column names are provided, whereas with table and column names, increased sparsity does not change the results much. This indicates that on enterprise data, LLMs rely primarily on the column names to predict the semantic types, maintaining almost constant accuracy even if no cell values are provided (Figure 4 left, sparsity of 1.0).

**Exp. 4: Improvement strategies.** Figure 3 shows that column type annotation accuracy decreases as the table width increases. To mitigate this effect, we explore three strategies for handling wide tables: *All columns (full table)* is our baseline setup, where the model predicts all semantic types of the full table in a single call. *Chunks of N columns* divides the table into chunks of  $N \in \{50, 10, 1\}$  columns and prompts the model for each chunk separately without the full table as context. *One column at a time (with full table context)* predicts each column type in a separate LLM call, but provides the full table in the prompt. We state the column to be predicted either by *name* or by *index* to analyze which variant works better.



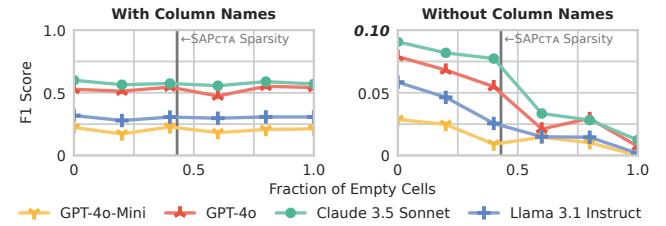
**Figure 3: Effect of the table widths.** The plots show support-weighted F1 scores for column type annotation with and without column names (zoomed in on the right). Increased table widths lead to worse results.

Table 5 shows that on public benchmarks, the differences between strategies are small, although smaller chunk sizes reduce performance if no column names are provided. This indicates that the data in public benchmarks provides strong contextual signals for column type annotation. For SAPCTA, however, chunking without column names does not further degrade the already low performance. With column names, chunking even improves results. Notably, for SAPCTA, the one-column-at-a-time strategy using the full table as context yields the best results. Since enterprise data without headers provides little signal, we suspect that the models rely primarily on the column names, largely ignoring the actual data. This method is also the most expensive, requiring one LLM call per column with the full table as context. Inference over the full SAPCTA dataset with GPT-4o costs USD 53.85, compared to USD 4.37 for the baseline—an order of magnitude more. Yet even with this setup, performance remains far below that on public benchmarks.

### 4.3 Discussion

**Learnings.** Our analysis of the customer data from SAP shows that enterprise data differs fundamentally from public benchmarks, and our experiments demonstrate that these differences lead to substantial performance decreases. The performance declines with higher sparsities and larger table widths, indicating that LLMs are unreliable at scale and sometimes fail to uphold even the basic table structure by generating an incorrect number of labels. Moreover, the non-descriptive schemas and complex data types pose significant challenges for LLMs. While we have shown this for the task of column type annotation, these data challenges will appear again throughout the following sections for other data engineering tasks, highlighting that they are general and affect a broad range of tasks.

**Outlook.** We believe that overcoming these challenges requires better representations for enterprise data. A first promising direction is the development of foundation models for relational data [60, 65]. While there is already a large body of research on foundation models for tabular data [11, 12, 17, 18, 33, 51, 66, 67], computing representations for complex enterprise data is still an open research problem. Some LLM weaknesses like their low reliability at scale require rethinking how they are applied to enterprise data, for example by using chunking to handle large tables. One way to alleviate the low self-expressiveness of the enterprise data could be to better contextualize the raw data by incorporating metadata such as



**Figure 4: Effect of the sparsity.** The plots show support-weighted F1 scores for column type annotation with and without column names (zoomed in on the right). Increased sparsity leads to worse results if no column names are given.

**Table 5: Strategies for handling wide tables, including our baseline setup predicting *all columns* in one call, *chunking*, and predicting just *one column with full table context* per call. The table shows support-weighted F1 scores for column type annotation with and without column names.**

Column Names	GitTablesCTA		SportsTables		SAPCTA	
	w/out	with	w/out	with	w/out	with
All columns (full table)	<b>0.52</b>	0.96	0.27	0.55	0.02	0.07
	<b>0.52</b>	0.96	<b>0.29</b>	0.51	0.01	0.09
	0.48	<b>0.97</b>	0.20	<b>0.62</b>	0.01	0.22
	0.22	0.92	0.06	0.61	0.01	0.20
1 column by index (full table)	0.30	0.66	0.07	0.11	<b>0.03</b>	0.11
1 column by name (full table)	-	0.89	-	0.50	-	<b>0.32</b>
All columns (full table)	<b>0.56</b>	<b>0.99</b>	0.57	<b>0.91</b>	0.03	0.27
	0.54	<b>0.99</b>	<b>0.58</b>	0.89	0.03	0.44
	0.51	<b>0.99</b>	0.45	<b>0.91</b>	0.03	0.46
	0.29	0.90	0.11	0.83	0.02	0.44
1 column by index (full table)	0.37	0.85	0.23	0.36	<b>0.05</b>	0.26
1 column by name (full table)	-	0.86	-	0.89	-	<b>0.52</b>

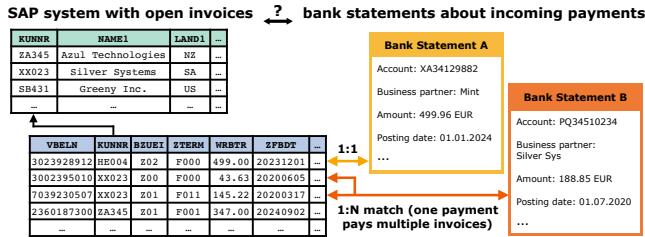
data dictionaries, which contain textual descriptions for table and column names and for symbolic values.

## 5 THE TASK CHALLENGE

A second dimension where data engineering in enterprise settings stands out is the complexity of the tasks. The academic definition of a data engineering task often makes simplifying assumptions that do not hold in real-world scenarios. For example, literature on entity matching often assumes that we compare the similarity of individual rows of two tables [28, 49]. By contrast, business entities in enterprises often span across multiple tables, forming graphs of tuples. At the same time, enterprises approach data engineering with broad business objectives in mind. Therefore, enterprise tasks are often composed of multiple steps (e.g., combining schema matching and entity matching). In this section, we discuss how enterprise tasks differ in their nature from the tasks evaluated in academia and examine how these differences affect LLMs.

### 5.1 Task Challenges

We base our analysis of the task-related challenges in enterprise data engineering on the following two case studies:



**Figure 5: Enterprise entity matching:** Bank statements of incoming payments must be matched to open invoices. Challenges include the invoices being represented by multiple tables, multi-match cases where a single payment pays multiple invoices, and discrepancies in amounts and descriptions.

**Case Study 1: Enterprise entity matching.** Entity matching is a data engineering task that often occurs in enterprise scenarios. We study the scenario shown in Figure 5, in which the bank statements of incoming payments must be matched to a company’s open invoices. The scenario showcases many challenges of entity matching in enterprises, including problems like complex table structures and the fact that matching requires finding 1:N or even N:M matches.

**Case Study 2: Enterprise database integration.** For our second case study, we look into another enterprise scenario where two customer databases from different companies must be merged. To ensure a unified and consistent customer database, this process involves the individual steps of schema matching, customer record matching, and data integration, as illustrated in Figure 2 (center).

In the following, we highlight the general challenges that arise in enterprise tasks and provide examples based on the two scenarios.

**C5: Entities span multiple tables.** Entities in enterprise systems are often business objects represented by multiple rows stored in different connected tables. In SAP systems, data pertaining to a particular material is scattered across the MARA (material type and basic statistics), MARC (manufacturing-related details), MBEW (valuation data), and other tables. The same holds true for the entities in our first case study (*entity matching*). For many data engineering tasks, one must therefore either manually construct views that extract the relevant fields into a single table, or approaches must work directly with the complex table structures that form a business entity.

**C6: Compound tasks.** Whereas data engineering in research is often addressed as isolated problems, such as deduplication [46] and missing value imputation [38], tasks in enterprise contexts are typically approached on a more holistic level concerning broader business objectives. Instead of focusing on individual tasks, enterprises aim to solve end-to-end workflows, as highlighted by our second case study (*database integration*). While the steps of such compound tasks could be executed sequentially, errors often propagate and amplify in later steps. Analyzing each step in isolation does not reveal the overall quality of the task in enterprise scenarios.

**C7: Task complexity.** Even individual workflow steps can be more complex than their counterparts studied in academia. For example, the entities to be matched in public entity matching datasets are usually of the same type, such as e-commerce products, restaurants, and

scholarly articles [28, 50]. By contrast, enterprise scenarios often require matching between different types of entities, like products to commodity codes or payments to invoices, which have overlapping but different sets of attributes. An additional challenge in enterprise scenarios is that the matches are often not 1:1 matches as in the literature, but can also be 1:N, N:1, or even N:M matches, making the problem much harder. In our payment-to-invoice matching scenario, it is common for a customer to pay multiple invoices with only one payment (1:N), or for one invoice to be paid by multiple payments (N:1), such as in the case of down payments.

**C8: Data complexity on top.** Along with these task complexities, the complexity of the data makes the tasks themselves also harder. For example, when matching payments to invoices, the memo lines of incoming bank statements are not standardized, but rather free-form text entered by customers. Therefore, complex errors occur regularly, making automated matching difficult. Compared to such data errors in enterprise scenarios, the errors in academic entity matching datasets are often much simpler.

## 5.2 Experiments & Results (Case Study 1)

To demonstrate how these challenges affect data engineering with LLMs, we first conduct experiments for our first case study (*entity matching*). Experiments for our second case study (*database integration*) follow in the next subsection.

**Setup.** For our first case study, we use a dataset of payments and invoices following the characteristics of the actual data from an enterprise system and reflecting the challenges described above (see Appendix A.2 in our extended technical report). The dataset contains 15,521 invoices and 12,332 payments. Following existing literature [41, 49], we formulate the entity matching task as a binary classification where the LLM decides if two entities (one payment and one invoice) match; i.e., if the payment pays for the invoice. Note that some of these combinations are part of multi-match cases where one payment pays for multiple invoices or vice versa. To keep the costs tractable, we use a subset of the dataset consisting of 790 payment-invoice combinations where the payment does pay for the invoice and 1,210 combinations where it does not.

For each combination, we prompt the LLM to predict if it is a match or a non-match. Furthermore, we experiment with presenting the invoices in two formats to the LLM: one format where each invoice is represented by multiple separate tables, and another format where each invoice is represented as one row in a flat view that combines all relevant columns from these tables, mimicking the traditional entity matching setting. We use a few-shot approach that includes one positive and one negative example in the prompt, as zero-shot performance was low on enterprise data. More examples did not improve the accuracy much. Starting from this simple setup, we make the task incrementally more complex.

**Exp. 5: Increasing task complexity.** In the following, we explain the different scenarios with results provided in Table 6:

1. *(Clean)* First, we evaluate on clean 1:1 matches using the flat table views described above, which is closest to the academic setting. In this setting, the payment’s memo line includes the correct reference numbers, the payment amount is exactly as stated in the invoice, and the customer name is also identical. As shown in Table 6 (*Clean*),

all models reach very high F1 scores in this setup. Next, we incrementally increase the difficulty based on the challenges we observe in the enterprise scenario.

2. (+ Errors) As described in C8, data errors in enterprises can differ from those in academic benchmarks, making the matching harder. To demonstrate this, we add representative errors and discrepancies that typically occur in real-world customer data to the payments and invoices, such as missing or additional digits in the reference numbers or minor discrepancies in the paid amount. As shown in Table 6 (+ Errors), these seemingly small inconsistencies already cause a noticeable drop in accuracy, highlighting the sensitivity of the models to such errors. We further study this effect in Exp. 6, showcasing that LLMs have substantial problems understanding the semantics of enterprise data.

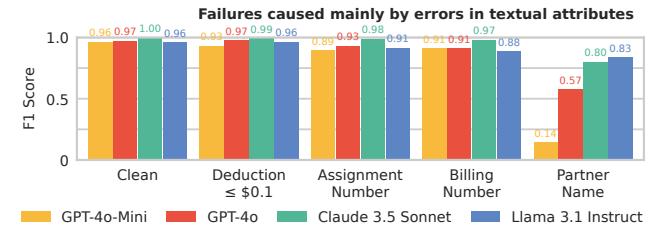
3. (+ Multi-Matches) Next, we make the matching even harder by focusing on the multi-match cases described in C7, where one payment pays for multiple invoices or multiple payments together pay for one invoice, a setting that is already closer to the real-world scenarios found in enterprises. As shown in Table 6 (+ Multi-Matches), this vastly increases the task’s difficulty, since amounts are split between multiple payments and payments might include multiple reference numbers for different invoices. To understand the root cause of the performance decrease and the high deviation between the models, we further investigate the precision and recall (see Appendix A.2 in our extended technical report). We find that while all models achieve a very high precision close to 1.00, the differences in F1 scores are primarily driven by variations in recall. This indicates that all models take a rather cautious approach, predicting matches only when they are highly certain. As a result, they prioritize high precision in matching and therefore miss many correct matches. While this high precision is sometimes preferred over high recall but low precision, the high rate of missed matches causes substantial manual efforts for enterprises to find the missing matches.

4. (+ Multiple Tables) Lastly, to show the impact of business entities spanning across multiple tables as described in C5, we represent invoices using multiple separate tables (as in the original SAP schema) instead of a single flat table view. While metadata about each invoice is stored in one table, specific information like the amount and due date is stored in a second table, and information about customers is stored in another separate customer table. In this scenario, which now closely resembles the actual enterprise task in its full complexity, all models see large performance decreases compared to the previous experiments, indicating that the models have difficulties working with the complex data structures used in enterprises.

**Exp. 6: Data errors and task complexity.** Our second experiment analyzes the impact of different error types on enterprise tasks like payment-to-invoice matching, highlighting typical data challenges that appear in enterprise scenarios. We insert individual types of data errors into the clean data from Exp. 5 to study their effects in isolation. We focus on four typical cases: (1) we deduct up to USD 0.1 from the paid amount, (2) we remove or change digits in the assignment or (3) in the billing number, and (4) we slightly vary the partner name, for example with abbreviations like *KL Technologies* instead of *Kim & Lee Technologies*. The results shown in Figure 6 indicate that even minor discrepancies can lead to performance

**Table 6: Increased task complexity when matching payments to invoices causes a steady decrease in F1 scores. The first column (*Clean*) resembles simple entity matching on data from public sources, while the last column (+ *Multiple Tables*) represents the enterprise scenario with all its complexities.**

	Clean	+ Errors	+ Multi-Matches	+ Multiple Tables
GPT-4o-Mini	0.98	0.58	0.53	0.45
GPT-4o	0.97	0.80	0.64	0.58
Claude 3.5 Sonnet	0.97	0.89	<b>0.86</b>	0.58
Llama 3.1 Instruct	<b>0.99</b>	<b>0.95</b>	0.81	<b>0.72</b>



**Figure 6: Impact of different error categories when matching payments to invoices. Discrepancies in the partner name attribute cause a substantially larger drop than numerical attributes, indicating a strong reliance on textual fields.**

degradations. Interestingly, errors in the numerical fields cause only modest performance declines, whereas discrepancies in the business partner names result in much more severe performance drops. This suggests that LLMs rely heavily on the availability of non-erroneous textual fields for entity matching, highlighting their inability to understand the complex numerical and symbolic values prominent in enterprise data.

### 5.3 Experiments & Results (Case Study 2)

To study the compound tasks described in C6, we now conduct experiments based on our second case study (*database integration*).

**Setup.** We use a dataset for the integration of two companies’ customer databases that is based on real-world data (see Appendix A.3 in our extended technical report). Company A’s database consists of a single table with 15 columns, selected as a subset of the SAP customer master data table KNA1 with columns such as KUNNR, NAME1, and ERDAT. Company B’s database consists of two tables connected by a foreign key relationship, with descriptive column names such as Organization Name, Address, and Tax Number. We experiment with datasets of different sizes, from 50 to 300 customers, to analyze the effects of the data scale. The customers in the two databases have an overlap of 60%, so 40% of the customers exist only in one of the databases. To evaluate the predicted result table, we compare it with the ground truth result table of integrated customers and compute the cell-level accuracy (i.e. the fraction of correct cells).

**Exp. 7: Error propagation in compound tasks.** In our first experiment, we analyze the effect of error propagation in compound

tasks with multiple steps. We formulate the database integration task as a sequence of three steps, as visualized in Figure 2 (center):

1. (*Schema Matching*) To integrate the data, we first perform schema matching to map the columns of Company B’s schema to the columns of Company A’s schema. Following existing research, we formulate the task as a binary classification between columns [41, 48]. Challenges arise because, for example, the address is represented by a single column at Company B, but split into multiple columns at Company A. We measure the accuracy of this step as the fraction of correctly found matches for the columns in Company A, since the output of the schema matching step is used to transform Company B’s database into Company A’s schema.

2. (*Entity Matching*) Next, we perform entity matching to determine which customers have entries in both companies’ databases. Similar to existing literature [41, 49], we formulate entity matching as a binary classification task: Given a pair of rows, one from Company A’s table and one from Company B’s transformed table, the LLM determines whether they refer to the same customer.

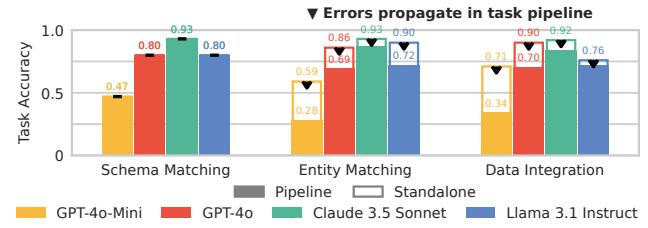
3. (*Data Integration*) To produce the integrated customer table, we use the LLM to merge the duplicate customer records found in both databases into a single row and to transform Company B’s remaining customer records into the format of Company A.

To measure the impact of error propagation when chaining the tasks, we run each step twice: once in a *pipeline* using the output of the previously executed step as input, and once in a *standalone* setting using the ground-truth output of the previous step as input. Figure 7 shows the results of this experiment. For the second (entity matching) and third (data integration) step of the pipeline, we observe that executing them in a pipeline leads to a decrease in accuracy compared to running them as standalone tasks. However, the decrease depends on the model, with GPT-4o-Mini seeing a more severe decrease than Claude 3.5 Sonnet and Llama 3.1 Instruct.

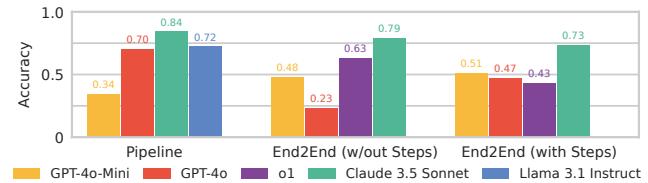
Interestingly, the accuracy between the entity matching and data integration steps in the pipeline does not decline and even slightly improves for some models. This shows a surprising effect for LLMs: While the models can make mistakes that propagate through the pipeline, they can, in some cases, also correct earlier mistakes. For example, some models fail to correctly associate the LAND1 field with the address during schema matching. However, when presented with three sample rows from Company A during data integration, they correctly populate the LAND1 field with the country mentioned in the address stored in Company B’s data.

Overall, these results suggest that error propagation can lead to a decrease in accuracy for compound tasks, but the significance depends on the LLM. In some cases, the effects are negligible.

**Exp. 8: End-to-end task execution.** A substantial overhead in enterprise scenarios stems from the fact that complex tasks must be manually decomposed into smaller steps. Recent LLMs, however, can use reasoning to solve multi-step problems on their own. One might now ask the question if the manual decomposition of the task is even needed, or if reasoning LLMs can break down the task on their own. To answer this question, we use the LLMs to directly integrate the two customer databases in a single LLM call and compare the results of this end-to-end execution to the pipelined execution from Exp. 7. We experiment with two prompting strategies: *w/o*



**Figure 7: Error propagation in our database integration case study with 100 customers.** Colored bars represent the *pipeline* setting, where errors carry over from previous steps, whereas stacked colorless bars show the *standalone* setting with correct inputs. The accuracy drop in the pipeline setting highlights the propagation of errors from earlier stages.



**Figure 8: Pipelined vs. end-to-end execution in our database integration case study with 100 customers.** Even at this small scale, end-to-end execution does not outperform pipeline execution.<sup>3</sup>



**Figure 9: Scaling up the number of customers for end-to-end database integration with OpenAI’s o1 drastically decreases the performance. For 300 customers, the model responded ‘Sorry, but I can’t fulfill that.’**

*steps* provides a textual description with instructions on how to perform the integration but does not mention how to break down the task, whereas *with steps* explicitly mentions the necessary steps (schema matching, entity matching, and data integration) in the prompt to provide a hint towards how to execute them.

Figure 8 shows the results of this experiment, including o1 from OpenAI as an additional LLM with reasoning capabilities. Most models perform better in the pipelined task execution, with only GPT-4o-Mini performing better in the end-to-end execution. When comparing the end-to-end execution *w/o steps* and *with steps*, we

<sup>3</sup>Three bars are missing from this plot: We did not execute the pipeline with o1 due to cost reasons, and Llama 3.1 Instruct did not produce responses for the end-to-end execution because of timeouts.

see mixed results. Whereas some LLMs benefit from describing the steps in the prompts, others do not.

**Exp. 9: End-to-end scaling.** While the end-to-end execution of compound tasks is attractive because it does not require manual task decomposition, it requires the full tables to be provided as input to the LLM because some steps like entity matching might need access to all tuples. This poses a problem, as larger tables lead to longer prompts that quickly exceed the LLM’s context window. As a last experiment in our database integration case study, we study the effects of scaling the number of customers in the company databases. Among the models evaluated in this work, only OpenAI’s o1 provides a sufficiently large output token limit (100K tokens) to integrate more than 100 customers in an end-to-end manner.

Figure 9 shows the results of scaling the total number of customers from 50 up to 300. We observe a drastic decline in performance as the number of customers increases, suggesting that the model struggles with long-range dependencies and finding matches in large datasets. When using 300 customers, the model even responds with “Sorry, but I can’t fulfill that”. Overall, LLMs remain limited in their ability to execute tasks at enterprise-scale in an end-to-end manner. In contrast, breaking up compound tasks into individual steps can, in principle, scale to larger tables, as tuples are processed one by one (e.g., by comparing all customers of databases A and B individually for entity matching).

## 5.4 Discussion

**Learnings.** Through our two case studies, we have demonstrated that automating enterprise tasks with LLMs remains extremely hard. First, we still need to rely on human effort since end-to-end execution with LLMs is brittle and does not scale well. This is consistent with Ashury-Tahan et al. [2], who found that even strong models often fail to perform robustly on complex table tasks. We also encountered problems in the output of some task formulations. For example, during end-to-end execution, where the output is a table with multiple rows, LLMs often generate varying numbers of columns for different rows.

Second, much of the complexity of enterprise tasks is inherent to business needs and cannot be avoided, such as having to deal with multi-match cases in entity matching. While reducing complexity by working with database views did improve results, constructing such views causes additional manual overheads since business objects in enterprises can be composed of tens or even hundreds of tables. Finally, some challenges like error propagation are not as severe as initially thought. Nevertheless, given their current abilities, LLMs are still not sufficient to achieve the level of performance required for enterprise-scale data engineering.

**Outlook.** We believe that LLMs out-of-the-box will not be sufficient to address the complexity of enterprise tasks because much of it is inherent to enterprise needs. Instead, complex tasks call for carefully-designed systems that use LLMs as reasoning or tool-calling agents while effectively incorporating human-LLM interactions. Wornow et al. [64] propose demonstrating workflows to foundation models, which is promising for automating frequently occurring tasks but still requires manual overheads. Another promising approach uses LLMs to create structured plans for compound tasks similar to recent advancements in LLMs for query planning [34, 59] before

executing these plans step-by-step using existing approaches (e.g., heuristics for schema matching) or even the LLM itself. This direction is reinforced by our findings: our experiments show that the stepwise execution of compound tasks improves performance over end-to-end execution while also enabling the LLM to correct earlier mistakes, for example by using information from subsequent steps.

## 6 THE KNOWLEDGE CHALLENGE

Since LLMs are primarily trained on public data, they lack information about a company’s internal business processes and policies, as well as about its proprietary tools and systems [26]. In this section, we analyze how the lack of company-specific knowledge in LLMs affects their accuracy on data engineering tasks.

### 6.1 Knowledge Challenges

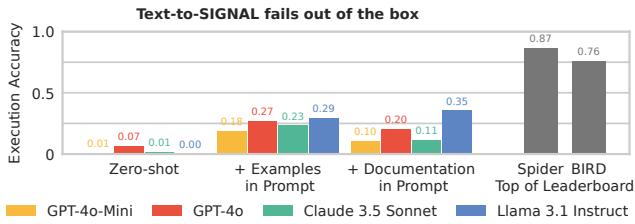
We identify two main categories of challenges: the lack of knowledge that is not available to LLMs but does exist in documentation and other company-internal sources, and company-specific extensions of databases that are typically not documented at all.

**C9: Proprietary but available knowledge.** Enterprises often use proprietary tools and systems to solve data engineering challenges. In contrast to well-established technologies like SQL, there is typically much less documentation about them available on the web. Therefore, it is reasonable to assume that LLMs trained on public data have little parametric knowledge about these tools and systems. One example is the domain-specific query language SIGNAL [23], which SAP provides to its customers for exploring data about business processes. SIGNAL resembles SQL in many aspects but also includes domain-specific features and syntax differences tailored towards process mining. While the language itself is well-specified on its public documentation page, it has a much smaller user base than SQL and, therefore, a smaller online footprint in terms of help pages, Q&A threads, and blog posts. Nevertheless, SIGNAL is frequently used by data engineers and must thus be well-supported by LLMs, for example, to translate natural language requests into SIGNAL queries.

**C10: Proprietary and unavailable knowledge.** Enterprise systems like those from SAP support company-specific changes and extensions to customize the system for individual customers. For example, customers can use hooks to add custom business logic. On the data level, customization means extending the database schema by adding customer-specific columns to existing tables or even additional tables to the customer namespace. In our analysis of real-world systems at SAP, we have come across thousands of such customer-defined tables. Since these changes are highly company-specific, they are typically not documented publicly—if they are documented at all. Moreover, customers sometimes “misuse” existing attributes of the standard SAP schema for different purposes. Such digressions from the public documentation pose significant challenges for data engineering with LLMs as well.

### 6.2 Experiments & Results

To analyze how challenge C9 affects data engineering with LLMs, we use LLMs to translate natural language requests into the domain-specific query language SIGNAL. The specifics of SIGNAL provide



**Figure 10: Text-to-SIGNAL execution accuracy.** The accuracy remains low even when including hand-picked examples and the SIGNAL documentation in the prompt. By contrast, LLMs achieve high accuracies on Text-to-SQL benchmarks [14, 15].

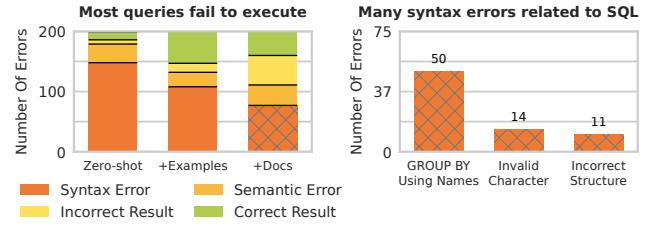
a good example for the proprietary knowledge primarily available in company-internal documentation. For challenge C10, we refer to Figure 1 as well as an additional experiment in Appendix A.5 of our extended technical report.

**Exp. 10: Text-to-SIGNAL.** In this experiment, we use LLMs to translate natural language requests into the proprietary query language SIGNAL, as shown in Figure 2 (right).

**Setup.** We experiment on a set of 200 randomly-sampled pairs of natural language requests and SIGNAL queries from a larger dataset in use at SAP. Given a request like *Retrieve the number of unique invoices from 'defaultview-255'*, we prompt the LLM to generate a corresponding SIGNAL query. We compare three approaches: First, we provide the LLM only with a short instruction explaining the task (*Zero-shot*). Next, we add three fixed example queries that were hand-picked from the SIGNAL documentation (*+ Examples*). Finally, we include both the examples as well as the documentation of the SIGNAL language in the prompt (*+ Documentation*). We manually tune our prompts in all scenarios to alleviate obvious mistakes and hint at the differences between SIGNAL and SQL. To evaluate the generated queries, we compute the execution accuracy by executing the ground truth and predicted query and comparing their results.

**Results.** Figure 10 (left) shows the results of this experiment. Only a small fraction of the generated SIGNAL queries are correct. Adding example queries to the prompt brings the biggest improvements in accuracy. By contrast, additionally including the documentation of the SIGNAL language in the prompt improves the accuracy only for some models, indicating that simply providing the required documentation during inference is not a viable solution. Finally, a comparison with the results on popular Text-to-SQL benchmarks [32, 68] (Figure 10 right) shows that the out-of-the-box accuracy for Text-to-SIGNAL is substantially lower.

**Why are examples and documentation not helping?** To better understand why exactly the models fail, we conduct an error analysis that categorizes the queries into those that execute with correct or incorrect results, and those that fail to execute because of syntactic or semantic errors. We focus on GPT-4o, the only model that correctly translated a non-negligible number of requests in all three configurations. Figure 11 (left) shows that most of the generated queries fail to execute because of syntactic or semantic errors, with the largest number of failures caused by syntax errors. Adding example queries and documentation to the prompt causes



**Figure 11: Text-to-SIGNAL errors for GPT-4o.** Most queries fail to execute because of syntactic or semantic errors. Examples and documentation in the prompt reduce syntax errors but cause semantic errors, as the model confuses parts of the documentation with the input question. Most syntax errors are caused by small differences between SIGNAL and SQL.

the number of syntax errors to decrease. However, the number of queries that contain semantic errors or return incorrect results increases as we add examples and documentation. A closer inspection of the generated queries reveals that the models often confuse parts of the documentation, like the column names used in explanations, with the request they have to translate. As a result, the number of fully correct queries is far from satisfactory even when including hand-picked examples and the documentation in the prompt.

**A bias towards SQL.** Another interesting effect is that the models tend to generate SQL-specific syntax and ignore the specifics of SIGNAL. To further dive into the root causes, we manually analyze the 75 syntax errors generated by GPT-4o in the (*+ Documentation*) setting and group them into three categories: the incorrect use of *GROUP BY* and *ORDER BY*, which in SIGNAL requires numerical indices instead of column names, the use of invalid characters, like an asterisk (\*) in a count statement which exists in SQL but not in SIGNAL, and the incorrect structure of the overall statement. As shown in Figure 11 (right), we find that even when explicitly including rules in the prompt to avoid such SQL-related syntax errors, the models still generate erroneous SIGNAL queries. This behavior indicates a strong bias of LLMs towards SQL, which may be caused by its much higher prominence in their training data. Similar problems could also exist in other scenarios where enterprise data and tasks differ only slightly from public knowledge, like currency conversion or company-specific data types.

### 6.3 Discussion

**Learnings.** Our experiments show that the lack of enterprise-specific knowledge affects the performance of LLMs. Even when including hand-picked examples and the full documentation in the prompt, LLMs are unable to actualize this information to solve enterprise tasks like Text-to-SIGNAL. An interesting finding here is that biases in the models’ knowledge towards public sources, like their understanding of SQL, are hard to overcome even with explicit instructions, making LLMs difficult to adapt without more heavyweight solutions like fine-tuning that cause additional overheads.

**Outlook.** Our results demonstrate that simple out-of-the-box approaches like few-shot prompting and including relevant documentation in the prompt are insufficient. While we believe that more

involved approaches, like fine-tuning LLMs for specific enterprise use cases, could help them acquire the necessary domain knowledge, these approaches require extensive work, for example to label the necessary data for fine-tuning for each individual task. Moreover, even self-supervised (task-agnostic) pre-training on enterprise corpora comes with important challenges: enterprise data often contains valuable and sensitive information, which raises significant privacy and copyright concerns. Therefore, we believe that future research should explore controlled alternatives, such as working with synthetic or anonymized corpora that reflect enterprise data characteristics [53]. Despite all these challenges, we believe that LLMs tuned to enterprise settings are a promising direction and ideas in that direction have recently gained traction in industry to adapt foundation models for enterprise environments.<sup>4</sup>

## 7 COSTS AT ENTERPRISE SCALE

The cost of using LLMs at enterprise scale is an important orthogonal dimension to the challenges discussed before. In this section, we want to briefly highlight three main cost drivers: the large scale of the data, the algorithmic complexity of the tasks, and the high tokenizer fertility [52], which is a general problem on tabular data.

**Data scale.** As described in Section 4, enterprise databases are substantially larger than typical data engineering benchmark datasets, with SAP databases often reaching multiple terabytes in size. Table 7 shows how this large size translates into high LLM costs. Processing one gigabyte of tabular data from our SAPCTA corpus with GPT-4o-Mini (OpenAI’s cheapest model) already costs about USD 71. Note that this cost does not include generating any output tokens, which are substantially more expensive than input tokens, as shown in Table 1. Nevertheless, the input token costs already provide a rough estimate of the expected high costs. For example, processing the entire customer database with GPT-4o-Mini would cost USD 463T, and using the o1 reasoning model would cost USD 46M. These high costs clearly prohibit using LLMs on large fractions of the enterprise data. This is especially problematic as our experiments in Sections 4-6 have shown that cheaper models like GPT-4o-Mini often underperform more expensive models.

**Algorithmic complexity.** As a second cost driver, we want to highlight the algorithmic complexity of the ways in which many data engineering tasks are currently approached with LLMs. One example is the task of entity matching, which current approaches address by using LLMs to compare pairs of entities [33, 49]. This approach has an algorithmic complexity of  $O(N \times M)$ , where  $N$  and  $M$  are the numbers of entities in each table, making it intractable even for medium-sized tables. For example, matching just 1,000 payments to 1,000 invoices already costs USD 1,462, and matching 10,000 payments to 10,000 invoices would cost USD 146,165. Future work should therefore focus on reducing the cost of LLM-based approaches by rethinking task formulations and combining LLMs with smaller fine-tuned models [37, 71] and other techniques. For example, for entity matching, Wang et al. [62] propose a task formulation where, given a row from one table, the LLM is asked to identify all matching rows from the other table in one step. This reduces the complexity to  $O(N)$  if the second table fits in context, and

<sup>4</sup><https://www.databricks.com/company/newsroom/press-releases/databricks-and-anthropic-sign-landmark-deal-bring-claude-models> (last accessed November 4, 2025)

**Table 7: Tokens per byte and cost per GB for textual and enterprise data. Enterprise data requires twice as many tokens per byte, leading to twice the cost per GB. Encoding the entire SAP customer database causes high costs for all models.**

Tokens per Byte	USD per GB		
	GPT-4o-Mini	GPT-4o	o1
Wikipedia	0.23	34	574
SAPCTA (CSV)	0.47	71	1,181
<b>USD for entire database:</b>	<b>462,923</b>	<b>7,715,380</b>	<b>46,292,282</b>

Pricing by OpenAI in February 2025.

to  $O(N \times M/B)$  when batching is used with batch size  $B$ , offering a substantial reduction in cost without compromising accuracy.

**Tokenizer fertility.** A final important cost driver is the high fertility of existing tokenizers on tabular data. The tokenizer fertility measures the number of tokens required to represent a single piece of text [52]. Since tokenizers are trained to represent natural language text, their vocabulary contains tokens even for long words. By contrast, tabular enterprise data contains lots of symbols and numerical values. As shown in Table 7, this results in twice the number of tokens required to represent each byte of enterprise data compared to natural language text from Wikipedia, leading to twice the cost per gigabyte. Even if the cost of new models continues to decrease, the costs for processing tabular data will thus remain twice as high. Therefore, future work should focus on creating new tokenizers specifically designed for enterprise data, as well as on creating more efficient textual representations for such data [58].

## 8 CONCLUSION

In this paper, we systematically analyzed the challenges involved in applying LLMs to real-world data engineering scenarios in enterprises. Through experiments with multiple LLMs on a diverse set of tasks, we have shown that the effectiveness of these models is substantially affected by enterprise-specific complexities along the dimensions of *data*, *tasks*, and *knowledge*. While the absolute numbers of the observed effects may change with more experiments and newer models, the large gap between enterprise and public scenarios shows the clear need for future research. On the same note, we see a strong need for enterprise-specific LLMs, with many recent activities emerging from research and industry. We hope that our insights and learnings provide a helpful guide and inspiration for future efforts to make LLMs viable for enterprise data engineering and support their adoption in industry.

## ACKNOWLEDGMENTS

This work has been supported by the BMBF and the state of Hesse as part of the NHR program and the HMWK cluster project 3AI. It was also partially funded by the LOEWE Spitzenprofessur of the state of Hesse and has benefited from early stages of funding from the Deutsche Forschungsgemeinschaft (DFG) under Germany’s Excellence Strategy (EXC-3057); funding will begin in 2026. We thank DFKI Darmstadt, SAP, and hessian.AI as well as Alexey Streltsov, Matthias Urban, and Furkan Karakocaoglu for their support.

## REFERENCES

- [1] Anthropic. 2024. Claude 3.5 Sonnet Model Card Addendum.
- [2] Shir Ashury-Tahan, Yifan Mai, Rajmohan C, et al. 2025. The Mighty ToRR: A Benchmark for Table Reasoning and Robustness. <https://doi.org/10.48550/arXiv.2502.19412> [cs].
- [3] Shraddha Barke, Christian Poelitz, Carina Negreanu, et al. 2024. Solving Data-centric Tasks using Large Language Models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 626–638. <https://doi.org/10.18653/v1/2024.findings-naacl.41>
- [4] Chandra Sekhar Bhagavatula, Thanapon Noraset, and Doug Downey. 2015. TabEL: Entity Linking in Web Tables. In *The Semantic Web - ISWC 2015*. Springer International Publishing, Cham, 425–441. [https://doi.org/10.1007/978-3-319-25007-6\\_25](https://doi.org/10.1007/978-3-319-25007-6_25)
- [5] Jan-Micha Bodensohn, Ulf Brackmann, Liane Vogel, Anupam Sanghi, and Carsten Binnig. 2024. Automating Enterprise Data Engineering with LLMs. In *NeurIPS 2024 Third Table Representation Learning Workshop*. <https://openreview.net/pdf?id=m85fYEjtDc>
- [6] Jan-Micha Bodensohn, Ulf Brackmann, Liane Vogel, Matthias Urban, Anupam Sanghi, and Carsten Binnig. 2024. LLMs for Data Engineering on Enterprise Data. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26–30, 2024*. VLDB.org. <https://vlbd.org/workshops/2024/proceedings/TaDA/TaDA.4.pdf>
- [7] Tom Brown, Benjamin Mann, Nick Ryder, et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6fbcb4967418fb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6fbcb4967418fb8ac142f64a-Paper.pdf)
- [8] Peter Baile Chen, Fabian Wenz, Yi Zhang, et al. 2024. BEAVER: An Enterprise Benchmark for Text-to-SQL. <https://doi.org/10.48550/arXiv.2409.02038> [cs].
- [9] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 16344–16359. [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/67d57c32e20fd0a7a302cb81d36e40d5-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/67d57c32e20fd0a7a302cb81d36e40d5-Paper-Conference.pdf)
- [10] DeepSeek-AI, Daya Guo, Dejian Yang, et al. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. [arXiv:2501.12948](https://arxiv.org/abs/2501.12948) [cs.CL]. <https://arxiv.org/abs/2501.12948>
- [11] Xiang Deng, Huan Sun, Alyssa Lees, You Wu, and Cong Yu. 2020. TURL: table understanding through representation learning. *Proc. VLDB Endow.* 14, 3 (Nov. 2020), 307–319. <https://doi.org/10.14778/3430915.3430921>
- [12] Julian Eisenschlos, Maharshi Gor, Thomas Müller, and William Cohen. 2021. MATE: Multi-view Attention for Table Transformer Efficiency. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 7606–7619. <https://doi.org/10.18653/v1/2021.emnlp-main.600>
- [13] Simon Frieder, Luca Pinchetti, Chevalier Chevalier, et al. 2023. Mathematical Capabilities of ChatGPT. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 27699–27744. [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/58168e8a92994655da3939e7cc0918-Paper-Datasets\\_and\\_Benchmarks.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/58168e8a92994655da3939e7cc0918-Paper-Datasets_and_Benchmarks.pdf)
- [14] Dawei Gao, Haibin Wang, Yaliang Li, et al. 2024. Text-to-SQL Empowered by Large Language Models: A Benchmark Evaluation. *Proc. VLDB Endow.* 17, 5 (Jan. 2024), 1132–1145. <https://doi.org/10.14778/3641204.3641221>
- [15] Yingqi Gao, Yifu Liu, Xiaoxia Li, et al. 2025. A Preview of XiYan-SQL: A Multi-Generator Ensemble Framework for Text-to-SQL. <https://doi.org/10.48550/arXiv.2411.08599> [cs].
- [16] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, et al. 2024. The Llama 3 Herd of Models. <https://doi.org/10.48550/arXiv.2407.21783> arXiv:2407.21783 [cs].
- [17] Xinyi He, Yihao Liu, Mengyu Zhou, et al. 2025. TableLoRA: Low-rank Adaptation on Table Structure Understanding for Large Language Models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2025, Vienna, Austria, July 27 - August 1, 2025, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, 22376–22391. <https://aclanthology.org/2025.acl-long.1090/>
- [18] Jonathan Herzig, Paweł Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. TaPas: Weakly Supervised Table Parsing via Pre-training. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 4320–4333. <https://doi.org/10.18653/v1/2020.acl-main.398>
- [19] Madelon Hulsebos, Çağatay Demiralp, and Paul Demiralp. 2021. GitTables benchmark - column type detection. <https://doi.org/10.5281/zenodo.5706316>
- [20] Madelon Hulsebos, Çağatay Demiralp, and Paul Groth. 2023. GitTables: A Large-Scale Corpus of Relational Tables. *Proceedings of the ACM on Management of Data* 1, 1 (May 2023), 30:1–30:17. <https://doi.org/10.1145/3588710>
- [21] Madelon Hulsebos, Kevin Hu, Michiel Bakker, et al. 2019. Sherlock: A Deep Learning Approach to Semantic Data Type Detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. Association for Computing Machinery, New York, NY, USA, 1500–1508. <https://doi.org/10.1145/3292500.3330993>
- [22] Gonzalo Jaimovich-López, César Ferri, José Hernández-Orallo, Fernando Martínez-Plumed, and María José Ramírez-Quintana. 2022. Can language models automate data wrangling? *Machine Learning* 112, 6 (2022), 2053–2082. <https://doi.org/10.1007/s10994-022-06259-9>
- [23] Timotheus Kampik, Andre Lücke, Jörn Horstmann, Mark Wheeler, and David Eickhoff. 2023. SIGNAL – The SAP Signavio Analytics Query Language. <https://doi.org/10.48550/arXiv.2304.06811> arXiv:2304.06811 [cs].
- [24] Jaewoo Kang and Jeffrey F. Naughton. 2003. On schema matching with opaque column names and data values. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data (SIGMOD '03)*. Association for Computing Machinery, New York, NY, USA, 205–216. <https://doi.org/10.1145/872757.872783>
- [25] Moe Kayali, Anton Lykov, Ilias Fountalis, Nikolaos Vasiloglou, Dan Olteanu, and Dan Suciu. 2024. Chorus: Foundation Models for Unified Data Discovery and Exploration. *Proc. VLDB Endow.* 17, 8 (April 2024), 2104–2114. <https://doi.org/10.14778/3659437.3659461>
- [26] Moe Kayali, Fabian Wenz, Nesime Tatbul, and Çağatay Demiralp. 2025. Mind the Data Gap: Bridging LLMs to Enterprise Data Integration. In *15th Annual Conference on Innovative Data Systems Research*.
- [27] Tassilo Klein, Clemens Biehl, Margarida Costa, Andre Sres, Jonas Kolk, and Johannes Hoffart. 2024. SALT: Sales Autocompletion Linked Business Tables Dataset. In *NeurIPS 2024 Third Table Representation Learning Workshop*. <https://openreview.net/forum?id=U2bELpkWr>
- [28] Pradap Konda, Sanjib Das, Paul Suganthan G. C., et al. 2016. Magellan: toward building entity matching management systems over data science stacks. *Proceedings of the VLDB Endowment* 9, 13 (Sept. 2016), 1581–1584. <https://doi.org/10.14778/3007263.3007314>
- [29] Keti Korini and Christian Bizer. 2023. Column Type Annotation using ChatGPT. In *Joint Proceedings of Workshops at the 49th International Conference on Very Large Data Bases (VLDB 2023), Vancouver, Canada, August 28 - September 1, 2023 (CEUR Workshop Proceedings)*, Vol. 3462. CEUR-WS.org. <https://ceur-ws.org/Vol-3462/TADA1.pdf>
- [30] Keti Korini, Ralph Peeters, and Christian Bizer. 2022. SOTAB: The WDC Schema.org Table Annotation Benchmark. *SemTab@ISWC 2022, October 23–27, 2022, Hangzhou, China (Virtual)* 3320 (2022), 14–19.
- [31] Sven Langenecker, Christoph Sturm, Christian Schalles, and Carsten Binnig. 2023. SportsTables: A new Corpus for Semantic Type Detection. (2023). <https://doi.org/10.18420/BTW2023-68> ISBN: 9783885797258 Publisher: Gesellschaft für Informatik e.V.
- [32] Jinyang Li, Binyuan Hui, Ge Qu, et al. 2023. Can LLM Already Serve as A Database Interface? A Big Bench for Large-Scale Database Grounded Text-to-SQLs. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). [http://papers.nips.cc/paper\\_files/paper/2023/hash/83fc8fab1710363050bbd1d4b8cc0021-Abstract-Datasets\\_and\\_Benchmarks.html](http://papers.nips.cc/paper_files/paper/2023/hash/83fc8fab1710363050bbd1d4b8cc0021-Abstract-Datasets_and_Benchmarks.html)
- [33] Peng Li, Yeye He, Dror Yashar, et al. 2024. Table-GPT: Table Fine-tuned GPT for Diverse Table Tasks. *Proc. ACM Manag. Data* 2, 3 (May 2024), 176:1–176:28. <https://doi.org/10.1145/3654979>
- [34] Chunwei Liu, Matthew Russo, Michael Cafarella, et al. 2025. Palimpsest: Optimizing AI-Powered Analytics with Declarative Query Processing. In *15th Annual Conference on Innovative Data Systems Research*.
- [35] Nelson F. Liu, Kevin Lin, John Hewitt, et al. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Trans. Assoc. Comput. Linguistics* 12 (2024), 157–173. [https://doi.org/10.1162/TACL\\_A\\_00638](https://doi.org/10.1162/TACL_A_00638)
- [36] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *Comput. Surveys* 55, 9 (Jan. 2023), 195:1–195:35. <https://doi.org/10.1145/3560815>
- [37] Yurong Liu, Eduardo Peña, Aécio S. R. Santos, Eden Wu, and Juliana Freire. 2025. Magneto: Combining Small and Large Language Models for Schema Matching. *Proc. VLDB Endow.* 18, 8 (2025), 2681–2694. <https://www.vlbd.org/pvldb/vol18/p2681-freire.pdf>
- [38] Yinan Mei, Shaoxu Song, Chenguang Fang, Haifeng Yang, Jingyun Fang, and Jiang Long. 2021. Capturing Semantics for Imputation with Pre-trained Language Models. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, Chania, Greece, 61–72. <https://doi.org/10.1109/ICDE51399.2021.00013>
- [39] Meta. 2024. Llama 3.1 Model Card.
- [40] Sidharth Mudgal, Han Li, Theodoros Rekatsinas, et al. 2018. Deep learning for entity matching: A design space exploration. In *Proceedings of the 2018 international*

- conference on management of data. 19–34.
- [41] Avanika Narayan, Ines Chami, Laurel Orr, and Christopher Ré. 2022. Can Foundation Models Wrangle Your Data? *Proceedings of the VLDB Endowment* 16, 4 (Dec. 2022), 738–746. <https://doi.org/10.14778/3574245.3574258>
- [42] OpenAI. 2024. GPT-4o System Card.
- [43] OpenAI. 2024. Learning to Reason with LLMs. <https://openai.com/index/learning-to-reason-with-langs/>. Accessed 18-02-2025.
- [44] OpenAI. 2024. OpenAI o1 System Card.
- [45] OpenAI, Josh Achiam, Steven Adler, et al. 2024. GPT-4 Technical Report. <http://arxiv.org/abs/2303.08774> arXiv:2303.08774 [cs].
- [46] Thorsten Papenbrock, Arvid Heise, and Felix Naumann. 2015. Progressive Duplicate Detection. *IEEE Transactions on Knowledge and Data Engineering* 27, 5 (May 2015), 1316–1329. <https://doi.org/10.1109/TKDE.2014.235966>
- [47] Simone Papicchio, Paolo Papotti, and Luca Cagliero. 2025. QATCH: Automatic Evaluation of SQL-Centric Tasks on Proprietary Data. *ACM Trans. Intell. Syst. Technol.* 16, 2 (2025), 45:1–45:26. <https://doi.org/10.1145/3712704>
- [48] Marcel Parciale, Brecht Vandervoort, Frank Neven, Liesbet M. Peeters, and Stijn Vansumeren. 2024. Schema Matching with Large Language Models: an Experimental Study. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26–30, 2024*. VLDB.org. <https://vlbd.org/workshops/2024/proceedings/TaDA/TaDA.8.pdf>
- [49] Ralph Peeters and Christian Bizer. 2023. Using ChatGPT for Entity Matching. In *New Trends in Databases and Information Systems - ADBIS 2023 Short Papers, Doctoral Consortium and Workshops: AIDMA, DOING, K-Gals, MADEISD, PeRS, Barcelona, Spain, September 4-7, 2023, Proceedings (Communications in Computer and Information Science)*, Vol. 1850. Springer, 221–230. [https://doi.org/10.1007/978-3-031-42941-5\\_20](https://doi.org/10.1007/978-3-031-42941-5_20)
- [50] Ralph Peeters, Reng Chiz Der, and Christian Bizer. 2024. WDC Products: A Multi-Dimensional Entity Matching Benchmark. In *Proceedings 27th International Conference on Extending Database Technology, EDBT 2024, Paestum, Italy, March 25 - March 28*, Letizia Tanca, Qiong Luo, Giuseppe Polese, Loredana Caruccio, Xavier Oriol, and Donatella Firmani (Eds.). OpenProceedings.org, 22–33. <https://doi.org/10.48786/EDBT.2024.03>
- [51] Jingang Qu, David Holzmüller, Gaël Varoquaux, and Marine Le Morvan. 2025. TabICL: A Tabular Foundation Model for In-Context Learning on Large Data. <https://doi.org/10.48550/arXiv.2502.05564> [cs].
- [52] Phillip Rust, Jonas Pfeiffer, Ivan Vulic, Sebastian Ruder, and Iryna Gurevych. 2021. How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 3118–3135. <https://doi.org/10.18653/v1/2021.acl-long.243>
- [53] Anupam Sanghi and Jayant R. Haritsa. 2023. Synthetic Data Generation for Enterprise DBMS. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 3585–3588. <https://doi.org/10.1109/ICDE55515.2023.000274>
- [54] Alexandra Savelieva, Andreas Mueller, Avrilia Floratou, et al. 2022. The Need for Tabular Representation Learning: An Industry Perspective. *Table Representation Learning Workshop at NeurIPS 2022* (2022).
- [55] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, et al. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- [56] Ananya Singha, José Cambrónero, and Sumit Gulwani. 2023. Tabular Representation, Noisy Operators, and Impacts on Table Structure Understanding Tasks in LLMs. *Table Representation Learning Workshop at NeurIPS 2023* (2023).
- [57] Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024. Table Meets LLM: Can Large Language Models Understand Structured Table Data? A Benchmark and Empirical Study. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. ACM, Merida Mexico, 645–654. <https://doi.org/10.1145/3616855.3635752>
- [58] Immanuel Trummer. 2024. Generating Succinct Descriptions of Database Schemas for Cost-Efficient Prompting of Large Language Models. *Proc. VLDB Endow.* 17, 11 (Aug. 2024), 3511–3523. <https://doi.org/10.14778/3681954.3682017>
- [59] Matthias Urban and Carsten Binnig. 2024. CAESURA: Language Models as Multi-Modal Query Planners. In *14th Conference on Innovative Data Systems Research, CIDR 2024, Chaminade, HI, USA, January 14–17, 2024*. www.cidrdb.org. <https://www.cidrdb.org/cidr2024/papers/p14-urban.pdf>
- [60] Liane Vogel, Benjamin Hilprecht, and Carsten Binnig. 2022. Towards Foundation Models for Relational Databases [Vision Paper]. In *NeurIPS 2022 First Table Representation Workshop*. [https://openreview.net/forum?id=s1KINOQq71\\_](https://openreview.net/forum?id=s1KINOQq71_)
- [61] Adrian Vogelgesang, Michael Haubenschild, Jan Finis, et al. 2018. Get Real: How Benchmarks Fail to Represent the Real World. In *Proceedings of the Workshop on Testing Database Systems*. ACM, Houston TX USA, 1–6. <https://doi.org/10.1145/3209950.3209952>
- [62] Tianshu Wang, Xiaoyang Chen, Hongyu Lin, et al. 2025. Match, Compare, or Select? An Investigation of Large Language Models for Entity Matching. In *Proceedings of the 31st International Conference on Computational Linguistics*, Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (Eds.). Association for Computational Linguistics, Abu Dhabi, UAE, 96–109. <https://aclanthology.org/2025.coling-main.8/>
- [63] Jason Wei, Maarten Bosma, Vincent Y. Zhao, et al. 2022. Finetuned Language Models are Zero-Shot Learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25–29, 2022*. OpenReview.net. <https://openreview.net/forum?id=gEZrGC0zdqR>
- [64] Michael Wornow, Avanika Narayan, Krista Opsahl-Ong, Quinn McIntyre, Nigam Shah, and Christopher Ré. 2024. Automating the Enterprise with Foundation Models. *Proc. VLDB Endow.* 17, 11 (2024), 2805–2812. <https://doi.org/10.14778/3681954.3681964>
- [65] Zhaomin Wu, Shida Wang, Ziyang Wang, and Bingsheng He. 2025. Learning Relational Tabular Data without Shared Features. <https://doi.org/10.48550/arXiv.2502.10125> arXiv:2502.10125 [cs].
- [66] Jingfeng Yang, Aditya Gupta, Shyam Upadhyay, Luheng He, Rahul Goel, and Shachi Paul. 2022. TableFormer: Robust Transformer Modeling for Table-Text Encoding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 528–537. <https://doi.org/10.18653/v1/2022.acl-long.40>
- [67] Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. TaBERT: Pretraining for Joint Understanding of Textual and Tabular Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8413–8426. <https://doi.org/10.18653/v1/2020.acl-main.745>
- [68] Tao Yu, Rui Zhang, Kai Yang, et al. 2018. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, 3911–3921. <https://doi.org/10.18653/v1/D18-1425>
- [69] Dan Zhang, Madelon Hulsebos, Yoshihiko Suhara, Çağatay Demiralp, Jinfeng Li, and Wang-Chiew Tan. 2020. Sato: contextual semantic type detection in tables. *Proceedings of the VLDB Endowment* 13, 12 (Aug. 2020), 1835–1848. <https://doi.org/10.14778/3407790.3407793>
- [70] Haochen Zhang, Yuyang Dong, Chuan Xiao, and Masafumi Oyamada. 2024. Large Language Models as Data Preprocessors. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26–30, 2024*. VLDB.org. <https://vlbd.org/workshops/2024/proceedings/TaDA/TaDA.11.pdf>
- [71] Zeyu Zhang, Paul Groth, Iacer Calixto, and Sebastian Schelter. 2024. Directions Towards Efficient and Automated Data Wrangling with Large Language Models. In *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*, 301–304. <https://doi.org/10.1109/ICDEW61823.2024.000044> ISSN: 2473-3490.
- [72] Erkang Zhu, Dong Deng, Fatemeh Nargesian, and Renée J. Miller. 2019. JOSIE: Overlap Set Similarity Search for Finding Joinable Tables in Data Lakes. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD '19)*. Association for Computing Machinery, New York, NY, USA, 847–864. <https://doi.org/10.1145/3299869.3300065>