



HAL
open science

Vectorizing string entries for data processing on tables: when are larger language models better?

Léo Grinsztajn, Myung Jun Kim, Edouard Oyallon, Gaël Varoquaux

► To cite this version:

Léo Grinsztajn, Myung Jun Kim, Edouard Oyallon, Gaël Varoquaux. Vectorizing string entries for data processing on tables: when are larger language models better?. 2023. hal-04345931

HAL Id: hal-04345931

<https://hal.science/hal-04345931v1>

Preprint submitted on 14 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Vectorizing string entries for data processing on tables: when are larger language models better?

1st Léo Grinsztajn

SODA

INRIA

leo.grinsztajn@inria.fr

2nd Myung Jun Kim

SODA

INRIA

3rd Edouard Oyallon

MLIA

CRNS, Sorbonne University

4th Gaël Varoquaux

SODA

INRIA

Abstract—There are increasingly efficient data processing pipelines that work on vectors of numbers, for instance most machine learning models, or vector databases for fast similarity search. These require converting the data to numbers. While this conversion is easy for simple numerical and categorical entries, databases are strife with text entries, such as names or descriptions. In the age of large language models, what’s the best strategies to vectorize tables entries, baring in mind that larger models entail more operational complexity? We study the benefits of language models in 14 analytical tasks on tables while varying the training size, as well as for a fuzzy join benchmark. We introduce a simple characterization of a column that reveals two settings: 1) a *dirty categories* setting, where strings share much similarities across entries, and conversely 2) a *diverse entries* setting. For dirty categories, pretrained language models bring little-to-no benefit compared to simpler string models. For diverse entries, we show that larger language models improve data processing. For these we investigate the complexity-performance tradeoffs and show that they reflect those of classic text embedding: larger models tend to perform better, but it is useful to fine tune them for embedding purposes.

Index Terms—tabular data, language models, data processing, join, data analytics

I. INTRODUCTION

While much of data engineering deals with discrete entries—categories, normalized entities, or open-ended text—there is a growing trend to use data representations made of numerical vectors. For instance, vector databases [1] use such representations in fast similarity searches for retrieval and fuzzy joins. Neural networks, which brought revolutions in many aspects of data processing, are also based on numerical vectors to represent the available information, including in natural language applications which deal solely with discrete tokens. However, for typical data tables, with columns containing entries of different nature and type, recent work has shown that bigger, more sophisticated, neural methods do not outperform simpler machine-learning models based on trees [2]. These tree-based methods handle discrete entries naturally, but struggle when the data cannot be represented as a moderate number of categories. In such a case, it is useful to combine them with representations of the string surface form of the entries [3].

Good vectorial representation of the string entries in tables remains crucial. Practitioners often rely on pretrained word embeddings developed in natural language processing [4] or numerical representations built from substrings [3]. Modern

natural language processing has moved on to much more elaborate architectures, using pretrained attentional architectures [5] which have evolved to large language models LLMs, such as LLaMa [6]. But vectorizing text with very large language models requires multiple expensive and rare high-end GPUs due to their memory footprint; it induces large energy consumption [7]. By contrast, table entries are typical fairly short strings. They seldom have the complex grammatical or narrative structures that pushed the development of language models of increasing depth and context window. This beg the question: what are the computational trade-off to create vectorial representations of string entries in tables? Are pretrained language models needed or are string representations enough? How complex should a model be? Given the cottage industry of language model—to date, the HuggingFace model hub has 42 000 models for text classification, 2 700 for sentence embedding—, which one to choose to embed text entries in tables? Evaluating many models for a given analysis is clearly impracticable; there is a dire need for guidelines.

Here we contribute a thorough empirical study of embedding of string entries in table for data processing. We consider two settings: 1) Data analytics, *ie* statistical analysis of records in a table, where we consider 14 supervised learning tasks, and 2) Data engineering, in particular table assembly, where we consider fuzzy-join: joining across 50 pairs of tables with imperfect alignment in the entity surface forms. We investigate more than 30 string embedding approaches. We show that a simple measure of the diversity across string entries enables separating columns on which string representations suffice, with entries that resemble “*dirty categories*”, and columns with more *diverse entries* on which large language models are beneficial. On the diverse entries, we show that the learnings from the text-embedding literature in natural language processing carry over to the data engineering settings.

Section II introduces the specific problem settings that we study and the related work on embedding entries. Section III then describes our benchmarking material: the datasets we use and the embedding methods that we survey. Finally, section IV details the results from the benchmark, highlighting various important trends, before we conclude in section IV-E, giving high-level recommendations to encode text entries for data processing.

II. CONTEXT AND RELATED WORKS

A. Problem setting: vectorization in data processing

Analytics: Analytical tasks on tables tackle, in general, estimation of statistical properties of the records (entries in a row). Often these properties are conditional estimates of one attribute in a row as a function of others; For instance, in a real-estate application, one might be interested in linking the expected price of properties to their features, such as age, number of rooms... Such estimations can be cast in a statistical learning framework [8]. The statistical estimation is formulated on a dataset of n observations $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$, where each observation consists of a feature vector $x_i \in \mathbb{R}^p$, the input attributes, and an outcome $y_i \in \mathbb{R}$ or $y_i \in 1, 2, \dots, K$, the target attribute. For practitioners, however, this setting typically only appear toward the end of a long data engineering process. First, text and categorical features must be vectorized, which is especially challenging for high-cardinality categorical features. Second, information is often distributed across multiple tables, and a time-consuming part of the data processing pipeline consists of carefully joining these different tables. This paper focuses on the text entries, which lead to significant challenges in the data processing operation. It explores a pipeline based on vectorizing these text entries prior to statistical learning or joining tables. A good embedding approach is one that makes downstream tasks –predictions, joins– more accurate.

Fuzzy join: Fuzzy join, –and the related similarity-join, fuzzy-matching, and entity resolution–, requires linking across different tables entries which refer to the same entity. We focus on the many-to-one join problem, where we want to enrich a base table with an auxiliary table (the reference table). More formally, as described in [9], if we denote L and R two input tables, where L serves as the reference table, a fuzzy join can be defined as a function $J : R \rightarrow L \cup \perp$, where \perp denote no match. Note that each element from R can match only one element in L , the reference table, while each elements of L can match many elements in R .

Fuzzy joining often makes use of Nearest Neighbor algorithms on a well chosen representation of the data. As for data analytics, we study a simple pipeline, were we vectorize text entries prior to using a Nearest Neighbor algorithm. A good embedding should make the downstream matching more accurate.

Vectorizing records: For both tasks, analytics with statistical learning and fuzzy joining, we investigate a simple tabular data processing pipeline: text and high-cardinality features are vectorized using a language model (and concatenated to the numerical features for tabular analytics) and fed into a classical machine learning model. While each ad-hoc module results from a complex learning process, their aggregation into a tabular data processing pipeline is straightforward.

Vectorizing can be applied offline, prior to data analysis, as it is computed row by row, and the resulting feature engineering can be reused across many analytical task. Such a reuse simplifies operations and decreases computational costs. But

it must be put in perspective with the operational costs of the embedding model.

B. Related work: many ways to represent table entries

Encoding high-cardinality features: Given a table with text entries, the traditional statistical literature often relies on One-Hot Encoding, but it falls short when dealing with high-cardinality categories, as it creates an explosion of the dimensionality of resulting embeddings. To alleviate the problem, various replacement methods have been suggested. Target Encoding is a competitive alternative that associates each category with the average value of the target variable [10], but it breaks when dealing with categories not seen during the training (out-of-vocabulary problem).

Character-level approaches based on substrings can generalize to unseen text and improve data processing tasks [3]. A central idea here is to count occurrences of sub strings, for instance defined by words or character-level n-grams. These counts can then be turned into low-dimensional embeddings with a matrix factorization, for instance a PCA after Tf-Idf renormalization (term frequency–inverse document frequency) to make the count distributions more suited for the square loss. A more advanced approach, yet fast and lightweight relies on MinHash sketching –a probabilistic approach to capturing Jaccard similarities between sub-string ensembles– to create embeddings that expose containment [3]. Sub-string level models are widely used as part of machine-learning software packages such as Scikit-Learn [11] or Skrub [12]. These approaches, however, can only rely on the regularity in the data, as they do not incorporate any outside semantic information.

Incorporating external information: Enhancing tabular data with external information, often referred to as feature enrichment, can significantly boost the prediction accuracy. If done manually, however, this process typically requires intensive labor from skilled data scientists, often involving painful joins and aggregations. To automate the process, Deep Feature Synthesis [13] greedily carries out joins and aggregations across tables. However, it is not applicable on large databases where it faces tractability challenges and results in extremely high-dimensional vectors.

To mitigate this issue, subsequent research has attempted to generate useful embeddings for entities within tabular data. [8] developed a method that learns embeddings from knowledge graphs. They demonstrated that such embeddings brings background information that enhances performance when incorporated into various tables. However, this approach requires a challenging step involving explicitly matching text entries between tables and knowledge graphs.

Language models for tabular data prediction: With the widespread use of language models, several works have been proposed to enhance predictions for tabular data. Given that they are trained on huge corpora of texts, the embeddings from the language models can provide useful background knowledge. For example, [14] observed that performance improved on one clinical dataset when using BERT-embeddings. Similarly, [3]

reported competitive results when employing this approach. Moreover, language models are robust to variations in text entries [15], which solves the issue of rigorous entity matching required when incorporating external information.

Additionally, several works extend the use of language models beyond embedding entities to enhance predictions. [16] leverages recent advancements in code generation with language models to automatically generate new features, retaining only those that boost performance. [17] and [18] directly fine-tune a language model on raw data, reporting good performance on very small datasets. These models rely both on the background knowledge and predictive abilities of language models, making it challenging to disentangle their respective contributions. In this work, we show how language models can bring in background information, as opposed to string models learned on the table at hand.

Probing: Starting with [19], researchers have been training simple models on intermediary activation of neural networks to uncover the information contained in these hidden states. The motivation of this line of work is often to better understand the inner workings of these models. In this paper, we use similar methods for a more practical aim: to easily extract vectorized information from textual entities. More closely related to our work, [20] shows that probing methods can extract detailed information about the spatial and temporal location of entities from large language models such as LLaMA2 [21].

Text embeddings: Sentence embeddings provide a compact way to represent a text and its information. For this reason, they are now used for various purposes, from text classification to paragraph retrieval.

While such embeddings can be directly extracted from language models pretrained on pretext tasks, [22] argues that the semantic information inside the model embeddings is not fully exploited without finetuning. This has led to a rich line of research on finetuning methods for sentence embeddings, using various methods such as contrastive training [23] [24] [25] [26], finetuning for classification on labeled sentence pairs datasets such as NLI or NQ [27] [26] or training to imitate slower but better-performing cross encoder models [28], which take a pair of sentence as input.

While these models and methods have been evaluated on various tasks [29], they have not been studied in the specific context of tabular data processing and analytics, where string entries are typically quite short and redundant, free-form text is scarce, and text embeddings are sometimes combined with numerical features.

Table models: Following pretrained-language models, the training scheme of these models have been tailored to inputs belonging to tables, leading to pretrained table models [30] [31]. Compared to their text-trained counterparts, these models have shown improved performances on table specific tasks such as row population, entity linking, or table fact verification. In this paper, we do not directly use models to solve table specific tasks, but rather attempt to vectorize table entries to improve performance on data analytics and preprocessing.

III. EXPERIMENTAL SETUP: PROBING ANALYTICS AND JOINS

A. An analytics benchmark: predicting an attribute value

To evaluate the performance of different text entries vectorization schemes for tabular analytics, we start by introducing a new classification benchmark on datasets containing both useful numerical features and text entries.

a) *Datasets:* We gathered datasets across multiple sources, mainly previous machine learning studies and kaggle competitions. Most machine-learning studies unfortunately focus on numerical data and we found 28 tabular datasets with at least one of the column being a text entry and with at least 1500 rows. Out these, 13 datasets (14 tasks) have at least one string column that is important for prediction ¹. The text features contained in these tables are diverse, as shown in table I.:

- 1) **Bikewale** [32] ² Information on bikes and scooters in India. The task is to predict the degree of price of automobiles.
- 2) **Clear Corpus** [33] ³: Generic information about the reading passage excerpts for elementary school students. The task is to predict the readability of the excerpts. The text feature is the name of the book, not the excerpt.
- 3) **Company Employees** ⁴: Information on companies with over 1,000 employees. The task is to predict the size range of the companies.
- 4) **Employee Salaries** ⁵: Information on salaries for employees of the Montgomery County, MD. The task is to predict the current annual salary range of the employees.
- 5) **Employee remuneration and expenses earning over 75000** ⁶ Remuneration and expenses for employees earning over \$75,000 per year. The task is to predict the remuneration of employees.
- 6) **Goodreads** [32] ⁷ Datasets containing information about books. The task is to predict the average rating of each book.
- 7) **Journal Influence:** Scientific journals and their descriptive features. The task is to predict the influence of a journal.
- 8) **Spotify** ⁸: Generic information on Spotify tracks with some associated audio features. The task is to predict the popularity of the albums.
- 9) **US Accidents** ⁹: Information of accidents in US cities between 2016 and 2020. From this dataset, two tasks are

¹On the 28 datasets we consider, 11 show ROC-AUC gains of less than 1% when including the text features, compared to using only the numerical features, and 14 show gains of less than 3%. These gains are computed by taking the biggest gains among OpenAI embeddings, Skrub MinHashEncoder, and the 3 best models in the MTEB benchmark. We restrict our analysis to the 14 datasets with gains greater than 3%.

²http://pages.cs.wisc.edu/~anhai/data/784_data/bikes/csv_files/bikewale.csv

³<https://www.commonlit.org/blog/introducing-the-clear-corpus-an-open-dataset-to-advance-research-28ff8cfea84a/>

⁴<https://www.kaggle.com/peopledatalabssf/free-7-million-company-dataset>

⁵<https://openml.org/d/42125>

⁶<https://opendata.vancouver.ca/explore/dataset/employee-remuneration-and-expenses-earning-over-75000/information/?disjunctive.department&disjunctive.title>

⁷http://pages.cs.wisc.edu/~anhai/data/784_data/books2/csv_files/goodreads.csv

⁸<https://www.kaggle.com/datasets/maharshipandya/spotify-tracks-dataset>

⁹https://smoosavi.org/datasets/us_accidents

Dataset	Column	Example	Ngrams / 1000 rows
Wine Review	Country	Portugal	294
Bikewale	Bike Name	Honda CB Twister Drum/Electric start	2878
Zomato	Location	Koramangala 1st Block	1121
Zomato	Name	Tandoor Garden	8491
Zomato	Dish Liked	Kaju Katli, Gulab Jamun, Petha	7595
Employee Salary	Department Name	Fire and Rescue Services	932
Spotify	Song name	She's So Mellow	9688
Spotify	Artist name	Brandtson	12605
Company Employees	Domain name	bajajfinserv.in	8936
Company Employees	Industry	food & beverages	2003
Journal Influence	Journal name	Acta Biomaterialia	8088
Goodreads	Description	Anarchist, journalist, drama critic, advocate of birth control and free love, Emma Goldman was the most famous-and notorious-woman in. . .	66423
Ramen Ratings	Brand	Sapporo Ichiban	3120
Ramen Ratings	Variety	Tom Yum Seafood Creamy	8975

TABLE I: Examples of text features in our datasets.

conducted: (1) the range of accident counts for the US cities (2) the severity of the reported accidents.

- 11) **US Presidential** [8]: Voting statistics in the 2020 US presidential election along with information on US counties. The task is to predict the range of voting numbers across US counties.
- 12) **Ramen ratings**¹⁰. The dataset contains ratings and characteristics of various ramens produced from multiple countries. The task is to predict the ratings of the ramens, by the Chicago Department of Buildings since 2006. The task is to predict the Total Fee.
- 13) **Wine reviews** [34]¹¹. The dataset contains wine ratings, as well as various such as price, winery or a small description. The task is to predict the rating.
- 14) **Zomato**¹². Information and reviews of restaurants in Bengaluru, India. The task is to predict the ratings of the restaurants.

b) *Text and numerical features processing*: We consider a feature to be a text feature if its cardinality is greater than thirty.

¹⁰<https://www.kaggle.com/datasets/residentmario/ramen-ratings>

¹¹<https://github.com/rogerioxavier/X-Wines>

¹²<https://www.kaggle.com/datasets/himanshupoddar/zomato-bangalore-restaurants>

Other features (low cardinality categorical, numerical features and datetime features) will be referred as "numerical features" for simplicity, and are vectorized independently. We use a OneHotEncoder for low cardinality variables, MinHashEncoder [3] for features with a cardinality greater than 10, and the DatetimeEncoder from the package Skrub [12] for datetime features (it transforms the datetime into features corresponding to the year, month, day, hour etc.). Numerical features are scaled with scikit-learn's [11] StandardScaler. Regression datasets are converted to binary classification, and all dataset are balanced. Except specified otherwise, we use sklearn's GradientBoostingClassifier as a classifier, as it is a strong baseline [2]. The same model is used on the text embeddings combined with numerical features. For a discussion of this choice, see IV-E.

c) *Evaluations*: We use the same sample size for all datasets. This size varies across experiments, as specified, and we limit ourselves to sample sizes below 5000, which still encompass a large part of the datasets used by practitioners [35]. Evaluations are always done on 7 cross-validation folds.

B. Entity resolution: A fuzzy join benchmark

We also investigate embeddings in the context of a common time-consuming data processing step: entity resolution. More precisely, we focus on the many-to-one fuzzy join problem, which is emblematic of situations where we aim to enrich a base table with auxiliary tables containing more detailed information, as described in II-A.

Datasets: For benchmarking, we take the 50 pairs of tables from [9]. These dataset pairs are constructed using multiple snapshot from Wikipedia, and using the natural variations in page names to get different names for similar entities.

Method: Our simple pipeline consists of using a 1-NearestNeighbor on vectorized representations of the rows. These representation are computed using language model embeddings, or using scikit-learn's TfidfVectorizer for comparison. We also compare this pipeline to AutoFuzzyJoin [9], a state-of-the-art unsupervised framework that can infer suitable fuzzy-join programs on given input tables. Note that for benchmarking, we use the datasets introduced in the AutoFuzzyJoin paper [9].

C. Text embedding methods surveyed

a) *Language models*: We aim to evaluate diverse language models. We first gathered models from the top of the MTEB benchmark¹³ [29]. In particular, we focus on the two models at the top for some experiments:

- BAAI's bge-large-en-v1.5 [36]¹⁴: a 335M parameters model pretrained on a large scale corpus, and finetuned on corpuses of text pairs.

¹³As indicated by this leaderboard around November 2023: <https://huggingface.co/spaces/mteb/leaderboard>

¹⁴<https://huggingface.co/BAAI/bge-large-en-v1.5>

- LLMrails’s ember-v1 ¹⁵: a 335M parameters model trained on an extensive corpus of text pairs.

We compare these models to OpenAI’s embeddings [25] through their API ¹⁶, using the model "text-embedding-ada-002" and to various non-finetuned models: pretrained-encoders Bert [5] and Roberta [37], pretrained decoders Mistral 7B-v0.1 [38], LLaMA 1 [6] and LLaMA 2 [21], as well as the Pythia models [39]. For these models, the embeddings are obtained by "mean pooling" except when specified otherwise [40], i.e they are obtained by averaging the embeddings of each token at the last layer of the model. To reduce the dimension of text embeddings, we use a PCA with 30 components if not specified otherwise. We study this choice in subsection IV-E and show that it is indeed a good default. Finally, we compare these models to the simpler word model Fasttext [41].

Table II lists all the specific models that we investigate, with their major characteristics.

b) Substring based approaches: For comparison, we also use character-level approach based on substrings. We use scikit-learn’s TfidfVectorizer ¹⁷ to create an embedding based on the occurrence of character-level ngrams. This embedding, which has the drawback of being very high-dimensional, is then handled like embeddings from language models. A more advanced and faster model we use is the MinHashEncoder [3], available through the package Skrub [12], which takes advantage of the min-hash approximation of the Jaccard to build encodings whose L_0 distances are approximations of the Jaccard of their ngrams sets. By default, we use 30 components to reach the same dimension that the reduced language model embeddings.

IV. RESULTS: GAUGING EMBEDDINGS FROM SIMPLE TO COMPLEX

A. Sophisticated string embeddings matter

We benchmark the performance of simple pipelines using entry embeddings in our two settings: prediction and many-to-one fuzzy join.

a) Prediction: Figure 1 shows the performance of two language model embedding methods (OpenAI’s ada-002 and BAAI’s BGE-large-en-v1.5) compared to using Skrub’s MinHashEncoder and sklearn’s TfidfVectorizer ¹⁸, two string-based models described in III-C. On average across our 14 analytic tasks and across all training sizes from 500 to 5000, more sophisticated embeddings improve task performance: the MinHashEncoder outperforms TF-IDF vectorization, and OpenAI’s text embedding is best. This order is preserved whether we consider only the text columns, or all columns for




Fig. 1: **Analytics: more sophisticated embedding improve performance** across varying training sizes using sklearn’s GradientBoostingClassifier. The ranks are computed across both settings (predicting from text + numerical entries and predicting only from text entries), but not across sample size, and averaged on 14 datasets.




Fig. 2: **Fuzzy join: pretrained language models improve upon string distances** Average rank on AutoFuzzyJoin many-to-one unsupervised join benchmark using different encoders. For each method except AutoFuzzyJoin, we use the 1-Nearest-Neighbor from Scikit-learn, adapting the fuzzy join implementation in Skrub. The best F1 score is taken for each dataset and method from 7 candidates by varying the precision / recall tradeoff parameter (threshold for Nearest-Neighbor, target for AutoFuzzyJoin) between 0.3 and 0.9

the analysis. Jointly modeling text and numerical columns brings a notable benefit, which underlines the benefit of representing text with vectors of numbers.

Table II shows the performance of all the models we evaluate, and in particular the mean difference with Skrub’s MinHashEncoder. We can see that, on average across the 14 tasks, all the language models that we investigate improve upon the MinHashEncoder.

b) Fuzzy Join: Figure 2 shows the performance of a simple approach: utilizing language model embeddings as input for a 1-Nearest-Neighbor algorithm. Using this simple pipeline with three strong language embedding models (see III-C), we show that this baseline outperforms the AutoFuzzyJoin algorithm [9], as well as a 1-Nearest-Neighbor using sklearn’s TfidfVectorizer, on the 50 datasets benchmark from [9] (see III-B).

¹⁵<https://huggingface.co/llmrails/ember-v1>

¹⁶accessed between October and December 2023

¹⁷which is equivalent to the perhaps better known CountVectorizer, followed by a TfidfTransformer

¹⁸for the TfidfVectorizer we only display the best set of parameters we found, which is using a ngram range of (2, 3) on characters. We varied the ngram range among {(1, 2), (1, 3), (2, 3), (2, 4)} both on characters and words, and with and without Tfidf transformation.

ROC-AUC gain from OpenAI embeddings over MinHashEncoder

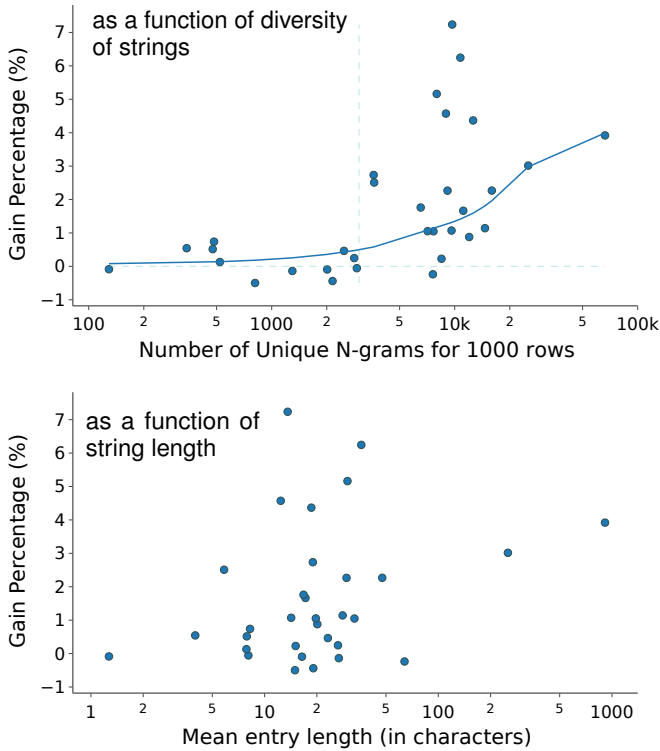


Fig. 3: **The number of unique ngrams per row predicts the gain better than the length of the text entries.** For every useful text column of 14 datasets, we compute the gain from replacing Skrub’s MinHashEncoder encoding of this column with OpenAI embeddings (while keeping the other columns encoded as before, see setup). The experiment is done at a train size of 1000. Each point corresponds to a text entry from one of our datasets. The ngrams are based on characters, and computed between lengths 2 and 4.

B. Two different regimes: dirty categories and diverse entries

Investigating the distribution of gains from using language model over substring-based methods reveals that the benefits are unevenly distributed. In Figure 3, we show the gain from using language model encodings over MinHashEncoder on each useful column¹⁹ belonging to the datasets in our tabular analytics benchmark. We see approximately zero gain for slightly less than half of the columns and significant gains for the other half. From the same Figure 3, we can separate the columns in two groups, based on a simple metric, the number of unique ngrams in the column for 1000 rows (computed on characters, between lengths of 2 and 4, for 1000 randomly sampled rows). This metric captures how the diversity of strings grows as a function of number of rows, revealing two regimes:

dirty categories columns where the number of unique ngrams is low, empirically below 3000 unique ngrams for 1000

¹⁹i.e where prediction is more than 0.5% better when including this column with either MinHashEncoder, OpenAI, or BAAI/bge-large-en-v1.5 embeddings over dropping it

ROC-AUC Gain for Embedding models over MinHash




Fig. 4: **Being better on classical embedding tasks translates to being better on tabular analytics.** But this is only the case in the *diverse entries* regime, where the number of unique ngrams in the column is large enough. The gain is computed for each column by replacing the MinHash encoding by a language model embedding (+ PCA), and averaged across columns (12 for ngrams < 3K, 20 for ngrams > 3K).

rows. On these columns, it seems that using a language model brings little benefits over string-based approaches. **diverse entries** columns where the number of unique ngrams is high, empirically above 3000 unique ngrams for 1000 rows. On these columns, using language model embeddings brings significant improvement.

Table I shows examples of columns belonging to these two categories. In contrast to our diversity metric, the length of the text entries has little relationship with the gain from using language model embeddings, as shown in Figure 3. Indeed a column may contain strings that are both very short, but also very diverse, as the *artist name* column of the *spotify* dataset, for which using OpenAI’s embedding over MinHashEncoder gives a ROC-AUC gain of 7.2%.

C. For diverse entries, using bigger, better models improves performance

The above shows that for *diverse entries*, using language model embeddings improves over simpler string-based methods. This begs the question, which embedding model should one use, among the enormous zoo of available models? Benchmarks such as MTEB [29] answer this questions for tasks like passage retrieval or sentiment analysis. Do the same tradeoffs apply to our case? Here text entries are much smaller than typical texts, and the resulting embeddings of string entries are combined with the other features of the tables before input to a subsequent machine-learning model.

a) *Comparison to embeddings benchmarks:* Nonetheless, Figure 4 shows that being better on the MTEB benchmark (on the average of the 56 tasks in the benchmark) quite directly translate to better performances on our tabular analytics tasks, in the *diverse entries* regime. In the *dirty categories* regime, in contrast, we see no gain from using better models.




Fig. 5: **Bigger is better, with diminishing returns** – We vary the model used to encode text features in place of Skrub’s MinHashEncoder. The gain percentage is averaged across 14 datasets, and computed on a train size of 1000, using a sklearn’s GradientBoostingClassifier. The error bars represent the standard error (halved for readability). Both text and numerical features are used. The models are taken from the top of the MTEB benchmark, except for the large decoder models (Pythia, Mistral, Llama) which are included to represent recent LLMs, and Bert, Roberta and DeBerta models. In a given model family bigger is better, but across all, very large models (Llama, sentence-t5, Pythia, mistral...) bring no significant benefits compared to e5-v2.




Fig. 6: **Being better on classical embedding tasks translates to being better for fuzzy joining.** The gain is averaged over 50 datasets, and computed by replacing the TfidfVectorizer by a language model embeddings before the Nearest Neighbor. Note that all the fuzzy join benchmark is in the "diverse entries" regime (more than 3000 unique ngrams for 1000 rows).

The fuzzy join benchmark described in III-B only contains columns in the *diverse entries* regime, with more than 3000 unique ngrams for 1000 rows. Quite logically, we also observe in Figure 6 that being better on the MTEB benchmark translates to being better on this benchmark as well.

b) *Bigger is better*: Embedding diverse entries of tables thus also follows the “bigger is better” scaling behavior described across a range of natural language tasks [42]. For a given family of models, figure 5 shows clear gains from increasing the model size. Existing pre-trained model families enable us to investing this trend for fine-tuned encoder models, such as e5 [26], but also decoder models with Pythia [39] models. For a given family we do not observe a plateau as we increase the model size.

In Table II, we also see that among the biggest models we evaluate, Mistral [38] and LLaMA 1 [6] and 2 [21] are on top of our leaderboard, despite being decoder models not finetuned for sentence similarity. This suggests that our pipeline will be able to benefit from both current and future advances in language models. This analysis could be extended to other

Model	Parameters	Model type	Fine tuned	ROC-AUC Gain (%)	Mean Rank (analytics)	F1 Gain (%)	MTEB score (Average)	MTEB score (Classification)
Llama-2-7b-hf	7.0B	Decoder	No	4.0	12.64	-30.0	Unknown	Unknown
Mistral-7B-v0.1	7.0B	Decoder	No	3.93	13.0	-20.6	Unknown	Unknown
e5-large-v2	335.1M	Encoder	Yes	3.77	10.0	2.6	62.25	75.24
llama-7b	7.0B	Decoder	No	3.75	14.0	-29.4	Unknown	Unknown
sentence-t5-xxl	4.9B	Encoder	Yes	3.31	17.29	-1.7	59.51	73.42
bge-large-en-v1.5	335.1M	Encoder	Yes	3.17	16.07	2.15	64.23	75.97
e5-large	335.1M	Encoder	Yes	3.09	14.57	2.5	61.42	73.14
OpenAI Ada-002	Unknown	Unknown	Yes	2.86	19.79	2.8	Unknown	Unknown
pythia-6.9b	6.9B	Decoder	No	2.81	21.21	-24.8	Unknown	Unknown
ember-v1	335.1M	Encoder	Yes	2.81	22.0	2.8	63.54	75.99
gte-large	335.1M	Encoder	Yes	2.74	20.0	3.8	63.13	73.33
gtr-t5-xxl	4.9B	Encoder	Yes	2.56	19.21	1.75	58.97	67.41
multilingual-e5-large	559.9M	Encoder	Yes	2.3	26.14	2.8	61.5	74.81
msmarco-bert-co-condensor	109.5M	Encoder	Yes	2.29	22.93	0.4	52.35	64.71
contriever-base-msmarco	109.5M	Encoder	Yes	1.96	29.36	1.5	56.0	66.68
jina-embedding-l-en-v1	334.9M	Encoder	Yes	1.96	30.36		Unknown	Unknown
roberta-base	125.0M	Encoder	No	1.74	31.5	-31.9	Unknown	Unknown
bert-base-cased	109.0M	Encoder	No	1.7	32.07	-20.0	Unknown	Unknown
all-MiniLM-L12-v2	33.4M	Encoder	Yes	1.54	30.29	-1.2	56.53	63.21
deberta-v3-large	335.0M	Encoder	No	1.53	36.86		Unknown	Unknown
Fasttext (cc-en)				1.53	38.64		Unknown	Unknown
all-distilroberta-v1	82.0M	Encoder	Yes	1.34	32.21	-1.7	Unknown	Unknown
bge-micro-v2	17.4M	Encoder	Yes	1.11	36.21	0.0	56.57	68.04
bge-micro	17.4M	Encoder	Yes	1.03	30.79	-0.1	55.71	66.35
paraphrase-multilingual-mpnet-base-v2	278.0M	Encoder	Yes	0.8	39.43	-3.5	Unknown	Unknown
paraphrase-multilingual-MiniLM-L12-v2	117.7M	Encoder	Yes	0.46	42.64	-9.3	Unknown	Unknown
Skrub MinHashEncoder				0.0	44.57		Unknown	Unknown

TABLE II: **Performances of various models averaged across 14 datasets** for analytics (ROC-AUC gain and mean rank), and for 50 datasets for fuzzy-join (F1 gain). Performances are computed for a sample size of 1000, and using our default pipeline (PCA with 30 components, GradientBoostingClassifier) for analytics. If a model comes as a suite of models, we only show the best performing one.

features known as being important for large language models, such as the training and finetuning data quantity. The worse performance of Pythia 6.9B is perhaps due to being trained on 300B tokens compared to 1 and 2T for LLaMA 1 and 2.

c) Finetuning: While a given model family exhibit a “bigger is better” scaling behavior on our tasks, finetuning the model for sentence embeddings is as important, maybe more. Indeed, in II and Figure 5, we see that small finetuned models like bge or e5 arrive at close or better performances than the largest models in our table while being an order of magnitude smaller (330M vs 7B parameters). Moreover, we see in Figure 5 that a better and newer finetuning procedure translates to bigger gain on the tabular analytics task, as can be seen comparing the different versions of a finetuned model like e5.

D. Language model can extract valuable knowledge from text features

We hypothesize that the performance gains from using language models to encode text entries come from the background knowledge contained in these models [20]. We provide some evidence for this claim in Figure 7, where the task is to

predict the population of European cities (with more than 10K inhabitants) from their name, and the names of their countries. Here, to ensure that the learner does not simply recognize the country of a city from its name –as city sizes differ between countries– the split between the train and test set is done using sklearn’s GroupKFold, such that the same country cannot appear both in the train and test set. We see that this makes it very hard for substring-based approach, as using Skrub’s MinHashEncoder leads to performance akin to random chance. On the contrary, using the OpenAI embedding, we are able to retain decent performances, suggesting that we are actually using the population knowledge contained inside the embedding.

E. A solid default pipeline

In this section, we check the robustness of our default pipeline using a series of ablations. The purpose is twofold: checking that the results of our experiments are not tainted by subpar settings, and guiding practitioners toward a simple yet effective pipeline. We recall that our default pipeline consists of encoding text entries with a language models, reducing the dimension of these embeddings with a Principal Component

R2 Score on European Cities' Population (log)




Fig. 7: Language models embed background knowledge – Comparison of the performance of Skrub’s MinHashEncoder and OpenAI’s embeddings for predicting the (log) population from the city and country names. This task is designed so that background knowledge is necessary for successful prediction: a city’s population cannot be concluded from its name. As typical city size vary across countries, we design the train and test splits so that cities from a same country appear either in the train or the test set, but not in both.

Analysis with 30 components, concatenating the results with the numerical features, and training a GradientBoostingClassifier on the result.

In Figure 8, we vary the number of components of the Principal Component Analysis used to reduce the embeddings dimension, and display the mean gain compared to using a dimension of 30, our default. We see that a dimension of 30 seems optimal until a sample size of 2000, and very close to optimal for bigger sample size until 5000 (the biggest size in our experiments).

In our paper, we kept the embedding dimension constant for all methods. In Figure 9, we vary the MinHashEncoder dimension while keeping the language model embedding dimension to 30 (using PCA). We see that up to 500, language models stay superior, with only 30 dimensions. We note that increasing the embedding dimension can leads to significantly higher downstream compute cost. Depending on how much embeddings can be reused, the higher cost of language model can be offset by using a smaller dimension.

Next, we study whether ensembling different models for text embeddings and numerical features beats our simple pipeline. Indeed, using a tree-based model on language model embeddings is unusual, and some work have shown that features are often linearly encoded in language models activations [20]. To this aim, we ensemble the prediction of a GradientBoostingClassifier trained on numerical features and a LogisticRegression trained on the text embeddings (without dimensionality reduction), and compute the mean ROC-AUC gain (across datasets) compared to our pipeline. The ensemble




Fig. 8: Choosing 30 components in the PCA is a reasonable choice. Comparing the performance of various choices of number of components for the PCA using OpenAI’s embeddings and evaluated using a GradientBoostingClassifier. While bigger number of components seems to improve slightly over 30 for bigger train sizes, 30 seems optimal for below 2000 train sizes, and close to optimal until 5000.

Using e5-large-v2 compared to MinHashEncoder, varying MinHashEncoder dimension




Fig. 9: Language models stay superior when increasing the MinHashEncoder dimension while keeping the language model embeddings to 30 dimension, though the gap narrows. We note that increasing this dimension leads to higher downstream compute cost.

bling is done either using scikit-learn’s VotingClassifier, i.e averaging the probability of each class, or using scikit-learn’s StackingClassifier, i.e training a LogisticRegression on the output of both ensemble models. As we can see in Figure 10, both embedding methods fail to improve upon our baseline on average. We do note however that on certain datasets, these methods bring improvements.




Fig. 10: Using a linear model on the embeddings underperforms our simple tree-based pipeline. ROC-AUC gain of using an ensemble of GradientBoostingClassifier on numerical features and a LogisticRegression on (OpenAI) embeddings (without dimensionality reduction), followed by either a stacking (using LogisticRegression) or (soft) voting using sklearn’s StackingClassifier or VotingClassifier, compared to our usual pipeline of using a PCA (dim 30) on embeddings, followed by a GradientBoostingClassifier on all features.

CONCLUSION

Rules of thumb: A thorough benchmark of embedding string entries for various data processing applications highlights trends precious for data engineering. These can be distilled in simple guidelines, good defaults to save practitioners time. First, it is useful to distinguish two kind of string columns: *dirty categories* with a low diversity across strings (for 1000 rows, no more than 3000 unique character-level n -grams with $n \in \{2, 3, 4\}$), and *diverse entries*. For dirty categories, lightweight string representations as the MinHashEncoder [3], [12] suffice. For diverse entries, borrowing language models from recent NLP developments brings much benefits. Here, bigger and more advanced language models to represent text entries in tables capture better knowledge useful for prediction and preprocessing on tables. For these columns, the findings from text embedding in natural language models carry over: larger models, fine-tuned to sentence-comparison tasks, bring benefits to analytic and entity resolution tasks. In particular, they markedly outperform word embeddings such as FastText which are currently often used as a default solution. Larger models come with increased computational burdens, and it can be useful to favor well fine-tuned models. To date, e5 (v2) [26] stands out as an excellent compromise.

Future work: Given a large database, better representations can be probably be obtained by adapting models to the database. However, this will increase markedly the computational and operational costs. The simple pipeline that we studied can easily be scaled to large datasets: the embedding complexity is linear with the number of records, and embeddings can be computed only once. Furthermore, progress in language model inference [43] [44] can make the embedding computation faster and cheaper. An interesting

avenue of research would be to study whether the particular background information we need for tabular analytics can be accessed without running the whole language model, as it has been observed that better information can be extracted from earlier layers in large language models [45].

REFERENCES

- [1] Y. Han, C. Liu, and P. Wang, “A Comprehensive Survey on Vector Database: Storage and Retrieval Technique, Challenge,” Oct. 2023.
- [2] L. Grinsztajn, E. Oyallon, and G. Varoquaux, “Why do tree-based models still outperform deep learning on tabular data?” Jul. 2022.
- [3] P. Cerda and G. Varoquaux, “Encoding high-cardinality string categorical variables,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 3, pp. 1164–1176, Mar. 2022.
- [4] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov, “Bag of Tricks for Efficient Text Classification,” in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, M. Lapata, P. Blunsom, and A. Koller, Eds. Valencia, Spain: Association for Computational Linguistics, Apr. 2017, pp. 427–431.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” May 2019.
- [6] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, “LLaMA: Open and Efficient Foundation Language Models,” Feb. 2023.
- [7] A. S. Luccioni, Y. Jernite, and E. Strubell, “Power Hungry Processing: Watts Driving the Cost of AI Deployment?” Nov. 2023.
- [8] A. Cvetkov-Iliev, A. Allauzen, and G. Varoquaux, “Relational Data Embeddings for Feature Enrichment with Background Information,” *Machine Learning*, vol. 112, 2022.
- [9] P. Li, X. Cheng, X. Chu, Y. He, and S. Chaudhuri, “Auto-FuzzyJoin: Auto-Program Fuzzy Similarity Joins Without Labeled Examples,” Mar. 2021.
- [10] D. Micci-Barreca, “A Preprocessing Scheme for High-Cardinality Categorical Attributes in Classification and Prediction Problems,” *SIGKDD Explorations*, vol. 3, pp. 27–32, Jul. 2001.
- [11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [12] S. Team, *Skrub: Prepping Tables for Machine Learning*, Inria Saclay, Palaiseau, France, 2023. [Online]. Available: <https://skrub-data.org/>
- [13] J. M. Kanter and K. Veeramachaneni, “Deep feature synthesis: Towards automating data science endeavors,” in *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. Campus des Cordeliers, Paris, France: IEEE, Oct. 2015, pp. 1–10.
- [14] K. V. Carballo, L. Na, Y. Ma, L. Boussioux, C. Zeng, L. R. Soenksen, and D. Bertsimas, “TabText: A Flexible and Contextual Approach to Tabular Data Representation,” Jul. 2023.
- [15] L. Chen, G. Varoquaux, and F. M. Suchanek, “Imputing out-of-vocabulary embeddings with LOVE makes language models robust with little cost,” *arXiv preprint arXiv:2203.07860*, 2022.
- [16] N. Hollmann, S. Müller, and F. Hutter, “Large Language Models for Automated Data Science: Introducing CAAFE for Context-Aware Automated Feature Engineering,” Sep. 2023.
- [17] S. Hegselmann, A. Buendia, H. Lang, M. Agrawal, X. Jiang, and D. Sonntag, “TabLLM: Few-shot Classification of Tabular Data with Large Language Models,” Mar. 2023.
- [18] T. Dinh, Y. Zeng, R. Zhang, Z. Lin, M. Gira, S. Rajput, J.-y. Sohn, D. Papailiopoulos, and K. Lee, “LIFT: Language-Interfaced Fine-Tuning for Non-language Machine Learning Tasks,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 11 763–11 784, Dec. 2022.
- [19] G. Alain and Y. Bengio, “Understanding intermediate layers using linear classifier probes,” Nov. 2018.
- [20] W. Gurnee and M. Tegmark, “Language Models Represent Space and Time,” Oct. 2023.

- [21] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom, "Llama 2: Open Foundation and Fine-Tuned Chat Models," Jul. 2023.
- [22] B. Li, H. Zhou, J. He, M. Wang, Y. Yang, and L. Li, "On the Sentence Embeddings from Pre-trained Language Models," Nov. 2020.
- [23] T. Gao, X. Yao, and D. Chen, "SimCSE: Simple Contrastive Learning of Sentence Embeddings," May 2022.
- [24] J. Ni, G. H. Abrego, N. Constant, J. Ma, K. B. Hall, D. Cer, and Y. Yang, "Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models," Dec. 2021.
- [25] A. Neelakantan, T. Xu, R. Puri, A. Radford, J. M. Han, J. Tworek, Q. Yuan, N. Tezak, J. W. Kim, C. Hallacy, J. Heidecke, P. Shyam, B. Power, T. E. Nekoul, G. Sastry, G. Krueger, D. Schnurr, F. P. Such, K. Hsu, M. Thompson, T. Khan, T. Sherbakov, J. Jang, P. Welinder, and L. Weng, "Text and Code Embeddings by Contrastive Pre-Training," Jan. 2022.
- [26] L. Wang, N. Yang, X. Huang, B. Jiao, L. Yang, D. Jiang, R. Majumder, and F. Wei, "Text Embeddings by Weakly-Supervised Contrastive Pre-training," Dec. 2022.
- [27] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, K. Toutanova, L. Jones, M. Kelcey, M.-W. Chang, A. M. Dai, J. Uszkoreit, Q. Le, and S. Petrov, "Natural Questions: A Benchmark for Question Answering Research," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 452–466, 2019.
- [28] N. Thakur, N. Reimers, J. Daxenberger, and I. Gurevych, "Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks," Apr. 2021.
- [29] N. Muennighoff, N. Tazi, L. Magne, and N. Reimers, "MTEB: Massive Text Embedding Benchmark," Mar. 2023.
- [30] X. Deng, H. Sun, A. Lees, Y. Wu, and C. Yu, "TURL: Table Understanding through Representation Learning," Dec. 2020.
- [31] T. Zhang, X. Yue, Y. Li, and H. Sun, "TableLlama: Towards Open Large Generalist Models for Tables," Nov. 2023.
- [32] G. C. P. S. G. C. K. P. G. Y. P. D. Das Sanjib, Doan AnHai, "The Magellan Data Repository."
- [33] S. Crossley, A. Heintz, J. S. Choi, J. Batchelor, M. Karimi, and A. Malatinszky, "A large-scaled corpus for assessing text readability," *Behavior Research Methods*, vol. 55, no. 2, pp. 491–507, 2023.
- [34] R. X. de Azambuja, A. J. Morais, and V. Filipe, "X-Wines: A Wine Dataset for Recommender Systems and Machine Learning," *Big Data and Cognitive Computing*, vol. 7, no. 1, p. 20, Mar. 2023.
- [35] "Largest Dataset Analyzed - Poll Results and Trends," <https://www.kdnuggets.com/largest-dataset-analyzed-poll-results-and-trends>.
- [36] S. Xiao, Z. Liu, P. Zhang, and N. Muennighof, "C-Pack: Packaged Resources To Advance General Chinese Embedding," Sep. 2023.
- [37] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," Jul. 2019.
- [38] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed, "Mistral 7B," Oct. 2023.
- [39] S. Biderman, H. Schoelkopf, Q. Anthony, H. Bradley, K. O'Brien, E. Hallahan, M. A. Khan, S. Purohit, U. S. Prashanth, E. Raff, A. Skowron, L. Sutawika, and O. van der Wal, "Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling," May 2023.
- [40] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," Aug. 2019.
- [41] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching Word Vectors with Subword Information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, Jun. 2017.
- [42] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling Laws for Neural Language Models," Jan. 2020.
- [43] F. Timbers, "Transformer inference tricks," <https://www.artfintel.com/p/transformer-inference-tricks>, Sep. 2023.
- [44] Y. Sheng, L. Zheng, B. Yuan, Z. Li, M. Ryabinin, D. Y. Fu, Z. Xie, B. Chen, C. Barrett, J. E. Gonzalez, P. Liang, C. Ré, I. Stoica, and C. Zhang, "FlexGen: High-Throughput Generative Inference of Large Language Models with a Single GPU," Jun. 2023.
- [45] K. Meng, D. Bau, A. Andonian, and Y. Belinkov, "Locating and Editing Factual Associations in GPT," Jan. 2023.