

Identifying the Units of Measurement in Tabular Data

Taha Ceritli^{1,2}[0000-0002-5059-8421] and
Christopher K. I. Williams^{1,2}[0000-0002-6270-4703]

¹ School of Informatics, University of Edinburgh, UK

² Alan Turing Institute, London, UK

Abstract. We consider the problem of identifying the units of measurement in a data column that contains both numeric values and unit symbols in each row, e.g., “5.2 l”, “7 pints”. In this case we seek to identify the dimension of the column (e.g. volume) and relate the unit symbols to valid units (e.g. litre, pint) obtained from a knowledge graph. Below we present *PUC*, a Probabilistic Unit Canonicalizer that can accurately identify the units of measurement, extract semantic descriptions of quantitative data columns and canonicalize their entries. We present the first messy real-world tabular datasets annotated for units of measurement, which can enable and accelerate the research in this area. Our experiments on these datasets show that PUC achieves better results than existing solutions.

Keywords: Data mining · Data understanding · Data cleaning.

1 Introduction

In a typical data analytics project, a large amount of time and effort is spent on understanding and cleaning the data. A particular task, which is often performed manually, is to understand how measurements are encoded in quantitative data columns. A *measurement* is a combination of a numeric value and a unit (e.g., litre), which can be encoded by several possible unit symbols for that unit (e.g., l and L).

Consider the following example: Suppose we are given the quantitative data column³ in Figure 1(a), which denotes the volume of freezers in various households. The value of each data entry is a measurement encoded by a unit symbol, except the last entry which does not have any unit symbol (we refer to the absence of a unit symbol in an entry as the *missing unit*). As shown in Figure 1(b), the measurements are encoded with two distinct units (litres and cubic feet) and six different unit symbols, some of which including `ltrs`, and `Cu` can be considered as *anomalous unit symbols* as they do not follow the standard encodings of

³ The data is sampled from the “Freezer_volume” column of the Household Electricity Survey (HES) dataset which can be accessed by registering at <https://tinyurl.com/ybbqu3n3>.

units. Moreover, a unit (e.g., litre) can be informative about the *dimension* of a measurement (in this case volume), which is defined as an expression of the dependence of a physical quantity on mutually independent components called *base quantities* in the International Vocabulary of Metrology⁴.

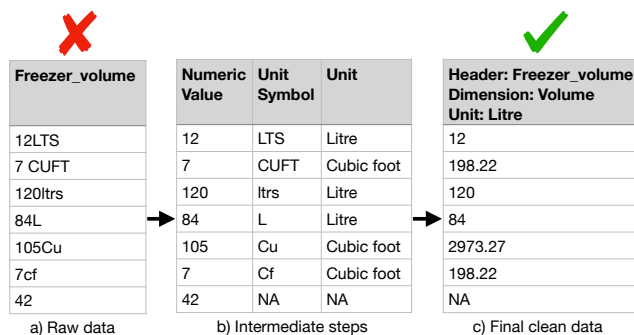


Fig. 1: A motivating real-world example that represents our pipeline. a) shows the samples of a raw dataset. b) indicates the intermediate steps required to transform the column. c) denotes the final data column obtained by applying the transformations.

We define *Unit Canonicalization (UC)* as the task of identifying the dimension of a data column and the units of all its entries so that the data entries can be canonicalized (i.e., so they all are expressed in the same units). See Figure 1(c) for an example. To address the UC problem, one needs to apply transformations such as parsing and identifying the units in the entries, inferring the common unit for the column which is placed in the metadata (e.g., the header), making the entries numeric and scaling the entries where needed. To the best of our knowledge, the UC problem is not addressed by any existing work in the literature (see Sec. 3 for a detailed discussion). In this work, we propose a probabilistic approach to the UC problem. Our contributions are as follows:

- We propose a probabilistic model which allows us to extract semantic information about a given data column containing unit symbols (such as its dimension and unit) and canonicalize its entries (Section 2).
- We make the first quantitative comparison of the existing methods on the unit identification task in real-world tabular data (Section 4).
- We present the first set of real-world datasets annotated for the units of measurement, to accelerate research in this area.

⁴ The document is accessible at https://www.bipm.org/documents/20126/2071204/JCGM_200_2012.pdf/f0e1ad45-d337-bbeb-53a6-15fe649d0ff1

2 Methodology

Representing Units: We represent units by extending a dictionary of units curated from Wikipedia⁵ with information from WikiData [19] and QUDT (Quantities, Units, Dimensions and Data Types Ontology)⁶. We associate each unit with a dimension and a list of unit symbols. See Appendix A for a detailed description.

The Proposed Model: We assume that a column of data $\mathbf{y} = \{y_i\}_{i=1}^N$ consisting of N rows has been read in, where each y_i denotes the characters in the i^{th} row. Each y_i is parsed using regular expressions (see Appendix C.2) to a numeric value v_i and a unit symbol x_i , which may be missing for some entries, i.e., x_i may be null. Our focus here is on the unit symbol x_i obtained. Denoting the number of possible dimensions for a column by K , our model has the following generative process:

$$\begin{aligned} \text{column dimension } t &\sim \mathcal{U}(1, K), \\ \text{row unit } u_i &\sim p(u_i|t), \\ \text{row label } z_i &= \begin{cases} u_i & \text{with probability } w_{u_i}^{u_i}, \\ m & \text{with probability } w_{u_i}^m, \\ a & \text{with probability } w_{u_i}^a, \end{cases} \\ \text{row symbol } x_i &\sim p(x_i|z_i), \end{aligned}$$

where \mathcal{U} denotes a discrete Uniform distribution, and $p(u_i|t)$ denotes how likely it is that a row unit u_i represents the dimension t . We model $p(u_i|t)$ with an indicator function that assigns a non-zero score only when a unit is a known unit of a dimension. Each row label z_i can be either the same as u_i , m or a according to the mixing proportions denoted by \mathbf{W} , where m and a respectively denote a missing or anomalous unit. Here the mixing proportions $w_{u_i}^{u_i} + w_{u_i}^m + w_{u_i}^a = 1$ for each row unit u_i . Since entries are often expected to be of a regular row label rather than the missing or anomalous labels, we favour regular labels during inference by using lower coefficients for the missing and anomalous labels, i.e. $w_{u_i}^m < w_{u_i}^{u_i}$ and $w_{u_i}^a < w_{u_i}^{u_i}$. These mixing proportions \mathbf{W} are assumed to be fixed and known. Even though one could also learn the mixing proportions, this may not be vital as long as the coefficients of the regular labels are larger than the others. Finally, $p(x_i|z_i)$ denotes our observation model which is explained below. Note that the above model is based on ptype [2], but extended to address unit canonicalization by incorporating semantic information about the units of measurements.

We build the observation model $p(x_i|z_i)$ upon three functions. First, we develop a Categorical distribution for row unit u_i where the categories correspond

⁵ Accessible at <https://github.com/marcolagi/quantulum/blob/master/quantulum/units.json>

⁶ <https://www.qudt.org/>

to the possible unit symbols for that unit, i.e., $p(x_i|z_i = u_i) = \text{Cat}(x_i, \pi_{u_i})$ where $\sum_{s=1}^{S_{u_i}} \pi_{u_i}^s = 1$. Here, $\pi_{u_i}^s$ denotes the probability of observing the s^{th} unit symbol for a unit u_i , while S_{u_i} denotes the number of known unit symbols for that unit. Secondly, we model missing units with an indicator function, which assigns a non-zero probability only when a unit symbol is missing. Lastly, we adapt the anomaly type in ptype, which is built based on the idea of an *X-factor* proposed by [14], to model anomalous unit symbols. Here, we introduce a likelihood function that assigns low but non-zero probabilities to any data value, which in turn allows the model to detect anomalous unit symbols which do not fit any known unit. Note that our model is fed with unit symbols, which are parsed using regular expressions. See Appendix C.2 for the details about our parsing step.

Inference: Given the row symbols \mathbf{x} in a data column, the initial task is to infer the column dimension t , which is cast to as the problem of calculating the posterior distribution of t given \mathbf{x} , namely $p(t|\mathbf{x})$. We then compute a posterior distribution over each row label conditioned on the dimension and the observed value, i.e., $p(z_i|t, x_i)$. Next, we determine the row units by calculating the posterior distribution of each row unit u_i given t , z_i and x_i , which is also used to predict the column unit. Lastly, we map a unit symbol labelled as anomalous to a known unit symbol and perform unit canonicalization. The derivations are presented for reproducibility in Appendix C.3.

3 Related Work

We are not aware of existing work specifically on the unit canonicalization problem. The closest related works are described below:

[18,9,16] use semantic web technologies to annotate quantitative data columns in terms of units. However, their annotations are based on ontological classes in domain-specific ontologies rather than the dimensions considered in this work. For example, [16] define the “High-Systolic-Blood-Pressure-Measurement” class with `kilopascal` as its unit. Given a data column and its ontological class, their goal is to represent all entries with the same predefined unit of the corresponding ontological class. [3] consider the task of annotating data entries in spreadsheets rather than tabular data. The authors split the label in a header (e.g., “Total Gallons”) into separate words (e.g., “Total” and “Gallons”), remove word inflections (e.g., obtaining “Gallon” from “Gallons”) and map word stems into known units and dimensions (e.g., “Gallon” and “volume”). Although their approach can be useful when information is explicitly given in a label, it would not work when the information is given implicitly, as in the label of “Credit Card Charges” which implicitly implies that the dimension is currency.

Regular expressions have been used for parsing and identifying units in text and unit conversion [8,21]. However, they are designed neither to annotate quantitative data columns with dimensions nor to canonicalize the units of their

entries. They do not even take the input as a data column, except for the professional version of Wolfram|Alpha which is not freely available. [17] develop a rule-based system named CCUT which maps unit symbols in tabular data to an ontology so that the data entries can be annotated with semantic information, which can be useful for table understanding. Unlike us, the authors do not use the contextual information in the entries of a data column, i.e., the unit symbols in the entries of a data column may be related to each other through the column dimension. Moreover, they do not consider the task of annotating quantitative data columns with dimensions.

Quantulum [11] use an ML model to disambiguate unit symbols in unstructured text (e.g., whether “pound” in a sentence refers to currency or mass). However, the entries of a data column consist of only numeric values and unit symbols, unlike long sentences where additional information is available through the other words. Similarly, [6] propose an ML framework named Grobid-Quantities (GQ) for processing text documents such as PDF files. However, identifying units in general tabular data can be more challenging since units in the scientific papers are more likely to follow standard notations. [20] apply simple ML techniques for dimension inference (e.g., the cosine similarities between word-embeddings representations of the header and the units), and then assign the most frequent unit for the inferred dimension as the column unit. Although this approach can be useful for dimension inference, it may be misleading for column unit inference as less frequent units can also be used to encode measurements in data columns. Lastly, dimension inference could benefit from Named Entity Recognition (NER) models. However, the set of tags supported by existing NER models are typically limited, e.g., Stanford NER includes only currency and time⁷.

4 Experiments

Experimental Setup: Here, we briefly describe the datasets, baseline methods and evaluation metrics used in our experiments. See Appendix B for a detailed description. The datasets, their sources and our annotations for dimensions and units can be accessed via <https://tinyurl.com/ay44c62r>. The code to reproduce the experiments are publicly available at <https://github.com/tahaceritli/puc>.

We conduct experiments on 24 data columns obtained from 16 CSV data files, each of which contains at least one column where the measurements are encoded by units of measurement. Their dimensions were annotated by hand for these sets, resulting in 2 currency, 2 data storage, 6 mass, 3 volume and 11 length columns. We also annotated each data entry in terms of its numeric value and unit symbol. We construct baselines for dimension inference by adapting CCUT, Grobid-Quantities (GQ), Pint, Stanford NER (S-NER) and Quantulum. On dimension inference, we compare our method with CCUT, GQ, Pint and Quantulum. To assess dimension inference, we use the overall accuracy and the Jaccard index. See [2] for a detailed description of these metrics and how they

⁷ See <https://nlp.stanford.edu/software/CRF-NER.shtml> for a complete list of the supported tags.

are used for type inference. Additionally, we evaluate the runtime of each method per data column (see Figure 3 in Supp. Mat.). To measure the performance on the unit identification task, we report the accuracies of the methods per dataset and apply paired t-tests to determine whether the predictions of the competitor methods are significantly different from the predictions of our method.

Results: Table 1 presents the performance of the methods on the column dimension inference task. The overall accuracies show that PUC performs better than the competitor methods. We observe a similar trend with the performance per dimension, quantified through the Jaccard index. These improvements are due to our model’s extensive knowledge about units, and its structure that takes into account the context shared among data rows. Note that the Jaccard index becomes zero when a method incorrectly labels all the data columns of a particular dimension (e.g., CCUT, GQ and Pint lead to zero Jaccard index for the currency dimension because they do not support currency symbols such as \$ and £ which occur in the entries of the “Currency” column of the Zomato dataset).

Dimension	Method					
	CCUT	GQ	Pint	S-NER	Quantulum	PUC
Currency	0.00	0.00	0.00	0.67	0.50	1.00
Data storage	0.00	0.00	0.00	0.00	0.50	1.00
Length	0.27	0.45	0.67	0.62	0.58	0.91
Mass	0.00	0.67	0.57	0.67	1.00	1.00
Volume	0.00	0.00	0.00	0.33	0.67	1.00
Overall Accuracy	0.12	0.38	0.50	0.71	0.71	0.96

Table 1: Performance of the methods for dimension inference using the Jaccard index and overall accuracy. We highlight the best score in each row by making the highest score bold.

Table 2 presents the accuracy of each method on each dataset (aggregated over columns) and the overall accuracy (averaging over all datasets). PUC performs consistently better than the competitor methods, although they are competitive with our method on some of the datasets. The performance of PUC on the Arabica and Zomato datasets are discussed in Appendix D. Note that accuracy becomes zero when a method fails to identify the unit of any data entry of a dataset (e.g., Pint cannot recognize the unit symbols in the MBA dataset such as LB and OZ which respectively denote the pound and ounce units). Finally, paired t-tests confirm that PUC is significantly different than each competitor method at the 0.05 level.

Dataset	CCUT	GQ	Pint	Quant.	PUC
Arabica	0.77	0.17	0.66	1.00	0.70
HES	0.18	0.00	0.27	0.64	0.98
Huffman	0.06	0.53	1.00	1.00	1.00
Maize	1.00	0.30	1.00	1.00	1.00
MBA	0.00	0.00	0.00	0.84	0.95
Open Units	0.00	0.00	1.00	0.00	1.00
PHM	0.99	0.95	0.99	1.00	1.00
query_2	0.00	0.00	1.00	0.00	1.00
query_4	0.00	0.00	1.00	0.00	1.00
Robusta	0.71	0.43	1.00	1.00	1.00
Zomato	0.00	0.00	0.00	0.00	0.60
143...23	0.00	0.69	0.00	0.00	0.97
143...62	0.58	0.00	0.00	0.95	0.95
228...96	1.00	0.00	1.00	1.00	1.00
3b5...ff	0.60	0.00	0.00	1.00	1.00
Overall Accuracy	0.39	0.20	0.59	0.62	0.94

Table 2: Accuracy of the methods on unit identification. We highlight the best score in each row by making the highest score bold. Quantulum is abbreviated as Quant.

References

1. Brown, G.: pyqudt (2019), <https://pypi.org/project/pyqudt/> [Accessed on 06/02/2020]
2. Ceritli, T., Williams, C.K.I., Geddes, J.: ptype: Probabilistic type inference. *Data Mining and Knowledge Discovery* **34**(3), 870–904 (2020)
3. Chambers, C., Erwig, M.: Reasoning about spreadsheets with labels and dimensions. *Journal of Visual Languages & Computing* **21**(5), 249–262 (2010)
4. Edwards, A.L.: Note on the "correction for continuity" in testing the significance of the difference between correlated proportions. *Psychometrika* **13**(3), 185–187 (1948)
5. Finkel, J.R., Grenager, T., Manning, C.: Incorporating non-local information into information extraction systems by Gibbs sampling. In: *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*. pp. 363–370. Association for Computational Linguistics (2005)
6. Foppiano, L., Romary, L., Ishii, M., Tanifuji, M.: Automatic identification and normalisation of physical measurements in scientific literature. In: *Proceedings of the ACM Symposium on Document Engineering 2019*. p. 24. ACM (2019)
7. Goldsmith, J.: wikipedia (2016), <https://pypi.org/project/wikipedia/> [Accessed on 06/02/2020]
8. Grecco, H.E.: pint Documentation Release 0.10.dev0 (2019), <https://buildmedia.readthedocs.org/media/pdf/pint/latest/pint.pdf> [Accessed on 05/08/2019]
9. Hignette, G., Buche, P., Dibie-Barthélemy, J., Haemmerlé, O.: Fuzzy annotation of web data tables driven by a domain ontology. In: *European Semantic Web Conference*. pp. 638–653. Springer Berlin Heidelberg (2009)
10. Keil, J.M., Schindler, S.: Comparison and evaluation of ontologies for units of measurement. *Semantic Web* **10**, 1–19 (08 2018). <https://doi.org/10.3233/SW-180310>
11. Lagi, M.: Quantulum (2016), <https://github.com/marcolagi/quantulum> [Accessed on 05/08/2019]
12. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. In: *Soviet Physics Doklady*. vol. 10, pp. 707–710 (1966)
13. Minhee, H.: Wikidata (2017), <https://pypi.org/project/wikipedia/> [Accessed on 06/02/2020]
14. Quinn, J.A., Williams, C.K.I., McIntosh, N.: Factorial switching linear dynamical systems applied to physiological condition monitoring. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(9), 1537–1551 (2009)
15. Rijgersberg, H., Van Assem, M., Top, J.: Ontology of units of measure and related concepts. *Semantic Web* **4**(1), 3–13 (2013)
16. Samadian, S., McManus, B., Wilkinson, M.D.: Automatic detection and resolution of measurement-unit conflicts in aggregated data. *BMC Medical Genomics* **7**(1), S12 (2014)
17. Shbita, B., Rajendran, A., Pujara, J., Knoblock, C.A.: Parsing, representing and transforming units of measure. *Modeling the World's Systems* (2019)
18. Van Assem, M., Rijgersberg, H., Wigham, M., Top, J.: Converting and annotating quantitative data tables. In: *International Semantic Web Conference*. pp. 16–31. Springer Berlin Heidelberg (2010)
19. Vrandečić, D., Krötzsch, M.: Wikidata: A free collaborative knowledgebase. *Communications of the ACM* **57**(10), 78–85 (09 2014). <https://doi.org/10.1145/2629489>, <http://doi.acm.org/10.1145/2629489>, <https://query.wikidata.org/> [Accessed on 25/11/2019]

20. Williams, J., Negreanu, C., Gordon, A.D., Sarkar, A.: Understanding and inferring units in spreadsheets. In: 2020 IEEE Symposium on Visual Languages and Human-Centric Computing. pp. 1–9 (2020)
21. Wolfram|Alpha: Units Overview-Wolfram Language Documentation (2019), <https://reference.wolfram.com/language/tutorial/UnitsOverview.html> [Accessed on 05/08/2019]

Supplemental Materials

Below, we give additional information about our knowledge graph of units (Appendix A), describe our experimental setup (Appendix B), provide additional information about our method (Appendix C) and discuss our experimental results further (Appendix D).

A Additional Information about Our Knowledge Graph

We represent units by extending a dictionary of units curated from Wikipedia⁸ with information from WikiData [19] and QUDT (Quantities, Units, Dimensions and Data Types Ontology)⁹. [10] show that WikiData is the most comprehensive knowledge graph for units and that QUDT contains additional information to WikiData. By extending the existing dictionary, we increase the number of units from 284 to 1080 (we only consider the units in English, although extensions to other languages are straightforward).

Table 3 presents two instances that respectively represent the units of litre and gram. The instances in the original dictionary can have six attributes: name, surfaces, entity, URI, dimensions, and symbols. Note that the “surfaces” attribute denotes a list of *strings* that refer to a unit, whereas the “symbols” attribute is a list of possible symbols and abbreviations for that unit.

name	surfaces	entity	URI	dimensions	symbols
litre	cubic decimetre, litre, volume	...	/wiki/Litre	{'base': 'decimetre', 1, L, ltr 'power': 3}	
gram	gram, gramme	mass	/wiki/Gram	—	g, gm

Table 3: Two elements of the unit dictionary. Note that URIs begin with <https://en.wikipedia.org/>.

For each instance we extract a dimension, a unit and a list of unit symbols. The name and entity attributes of the instances in Table 3 are respectively used as units and their dimensions. Note that the naming convention used in the dictionary differs from our terminology in that it refers to dimension as entity and uses the “dimensions” attribute to encode the relationship between units. The “surfaces” and “symbols” attributes are combined to build a set of unit symbols. To search for additional symbols of a unit, we query WikiData using the URI attribute. Additionally, the QUDT reference ID, when available in the response obtained from WikiData, is used to query QUDT.

⁸ Accessible at <https://github.com/marcolagi/quantulum/blob/master/quantulum/units.json>

⁹ <https://www.qudt.org/>

We support the following dimensions: acceleration, amount of substance, angle, area, capacitance, catalytic activity, charge, currency, current, data storage, data transfer rate, dimensionless, dynamic viscosity, electric potential, electrical conductance, electrical resistance, energy, flux density, force, frequency, illuminance, inductance, instance frequency, irradiance, kinematic viscosity, length, linear mass density, luminance, luminous flux, luminous intensity, magnetic field, magnetic flux, magnetomotive force, mass, mass flow, power, pressure, radiation absorbed dose, radiation exposure, radioactivity, sound level, speed, temperature, time, torque, typographical element, volume, volume (lumber), volumetric flow.

To query WikiData, Wikipedia and QUDT, we respectively use wikidata [13], wikipedia [7] and pyqudt [1].

B Additional Information about Our Experimental Setup

Datasets We conduct experiments on 24 data columns obtained from 16 CSV data files, each of which contains at least one column where the measurements are encoded by units of measurement. The datasets used can be briefly described as follows:

- Arabica, Robusta: a collection of reviews about coffee beans.
- HES (Household Electricity Survey): time series measurements of the electricity use of domestic appliances (to gain access to the data, please register at <https://tinyurl.com/ybbqu3n3>).
- Huffman: the Huffman Prairie flight trials in 1904, which is available through the U.S. Centennial of Flight Commission.
- Maize Meal: a list of maize meal products.
- MBA: a list of products in a grocery shop.
- Open Units: a list of 1,000 standard servings of branded drinks and their alcohol content.
- PHM Collection: information on the collection of Powerhouse Museum Sydney, including textual descriptions, physical, temporal, and spatial data as well as, where possible, thumbnail images.
- 143...6, 143...23 and 228...96: a set of data Web tables from T2Dv2 Gold Standard to evaluate matching systems on the task of matching Web tables to the DBpedia knowledge base.
- query_2, query_4: a set of tables extracted from WikiData using the properties of height or weight.
- Zomato: information about restaurants extracted from Zomato.

The dimensions of the data columns were annotated by hand for these sets, resulting in 2 currency, 2 data storage, 6 mass, 3 volume and 11 length columns. We also annotated each data entry in terms of its numeric value and unit symbol. Note that numeric values are missing in 5 data columns, i.e., only unit symbols are observed in the entries. For example, the "Quantity Units" column of the Open Units dataset consists of three unique values. These are `ml`, `pint` and `cl`,

which respectively denote the units of millilitres, pint and centilitres. Table 4 presents the number of entries, unique entries, units and unit symbols per data column. The number of unit symbols per data column varies between 2 and 11.

dataset	column	# non-missing entries	# unique entries	# units	# unit symbols
Arabica	Bag W...	1,283	51	2	2
Arabica	Altitude	132	81	2	11
HES	Freezer...	50	47	2	11
HES	Refrig...	38	37	2	9
Huffman	DIST...	71	68	2	2
Maize	PACK...	30	10	2	2
MBA	CURR...	43	19	3	6
Open U...	Quanti...	1,082	3	3	3
PHM	Height	30,312	1,313	3	5
PHM	Weight	179	135	2	3
PHM	Width	36,330	1,282	3	5
PHM	Depth	19,400	1,073	3	4
PHM	Diameter	2,106	377	2	2
Robusta	Bag W...	28	4	2	2
Robusta	Altitude	8	3	1	2
query_2	unitH...	22	3	3	3
query_2	unitW...	22	3	3	3
query_4	unitH...	22	3	3	3
Zomato	currency	6,386	9	9	9
143...62	FORMAT	123	19	2	3
143...23	Size	3,855	1,667	2	2
228...96	Size	377	256	2	2
3b5...ff	amount	5	5	2	2

Table 4: Size of the datasets used and the number of units and unit symbols in each data column.

Baselines As we describe in Sec. 3, the methods proposed in [18,9,16,17,6,8,11,5] do not address the unit canonicalization task. However, we can construct baselines to be used in our experiments by adapting some of these methods, namely CCUT [17], Grobid-Quantities (GQ) [6], Pint [8], Stanford NER (S-NER, [5]) and Quantum [11]. All these methods take as input a sentence (e.g., "...2 litres of water.") and annotate the words that refer to quantities (e.g., 2 litres) with their dimensions (e.g., volume), except S-NER which needs to be trained again for the dimension prediction. Here, we use the data values of the entries as inputs to these methods.

We construct baselines for dimension inference as follows. We first identify the dimensions of the entries of a data column using a competitor method and then predict the column dimension through majority voting, i.e., assigning the most common dimension as the column dimension. The baseline method based on the S-NER is obtained by training the pre-trained model with pairs of unit

symbols and their dimensions (e.g., **metres** - length). The resulting model then generates a tag for each data entry of a data column and can be used to assign a dimension to that column through majority voting.

On the unit identification task, we evaluate whether the unit of a data entry can be correctly identified, e.g., **1 cm** as **1** and **centimetre**. For this, we compare our method with CCUT, GQ, Pint and Quantum, which are developed for identifying units in textual documents.

Evaluation Metrics We use different sets of metrics for dimension inference and unit identification. For dimension inference, we use the overall accuracy and the Jaccard index. See [2] for a detailed description of these metrics and how they are used for type inference. To measure the performance on the unit identification task, we report the accuracies of the methods per dataset and apply paired t-tests to determine whether the predictions of the competitor methods are significantly different from the predictions of our method. Note that the accuracy on the unit identification task measures the ratio of number of entries in a dataset for which the unit is correctly predicted over the number of entries in a dataset.

C Additional Information about Our Method

C.1 Notation

Table 5 presents a summary of our notation.

Symbol	Description
t	the column dimension
u_i	the unit of the i^{th} row
z_i	the label of the i^{th} row
v_i	the numeric value of the i^{th} row
x_i	the unit symbol of the i^{th} row
y_i	the characters of the i^{th} row
K	the number of possible dimensions
L_t	the number of possible units for dimension t
S_{u_i}	the number of possible unit symbols for unit u_i

Table 5: A summary of the notation used by PUC.

We assume that a column of data $\mathbf{y} = \{y_i\}_{i=1}^N$ consisting of N rows has been read in, where each y_i denotes the characters in the i^{th} row. Additionally, each y_i is assumed to be parsed to a numeric value v_i and a unit symbol x_i , which may be missing for some entries, i.e., x_i may be null. We propose a generative model with a set of latent variables t , $\mathbf{u} = \{u_i\}_{i=1}^N$ and $\mathbf{z} = \{z_i\}_{i=1}^N$, where t denotes the dimension of a column, u_i the unit and z_i the label of its i^{th} row.

The missing and anomalous labels, denoted by m and a respectively, are used to model the data entries where the unit symbols are missing or anomalous. Thus, each z_i can be m or a as well as a row unit that fit the column dimension, i.e. $z_i \in \{\text{Litre, Cubic foot, ..., } m, a\}$ given that t is volume. With this noise model, we make our inference procedure robust against missing and anomalous unit symbols.

C.2 Parsing Unit Symbols

Unit symbols are usually positioned after quantities as in 1 L, with some exceptions where the conventions are different. For example, they are usually placed before quantities to represent monetary amounts, e.g., \$159000 and \$85810. When abbreviations are used, however, unit symbols are placed after numeric parts, e.g., 70 USD, 19.68 AUD. Monetary amounts can also be represented in various non-standard formats. For example, whitespace may be placed between symbols and amounts, e.g., \$ 1012, \$ 964. We develop regular expressions by taking into account possible positions of unit symbols. In particular, we can parse a numeric value expressed as a fraction, or an integer or a float, followed by whitespace, then a string. Additionally, we remove leading and trailing whitespace as well as trailing dots. We use the following regular expression to parse a unit symbol from a given text:

```
import re
numeric_with_string_const_pattern = r"""
    [-+]? # optional sign
    (
        (?: \d+ \/ \d+ ) # 1/4 etc
        |
        (?: \d* [.,] \d+ ) # .1 .12 .123 etc 9.1 etc 98.1 etc
        |
        (?: \d+ \.? ) # 1. 12. 123. etc 1 12 123 etc
    ) ?
    # whitespace as separator
    (?: [\s]*) ?

    # followed by optional characters (alphanumeric, whitespace, some punctuation marks
    ([\w\s.!?\-\-]*) ?
    """
rx = re.compile(numeric_with_string_const_pattern, re.VERBOSE)
```

C.3 Derivations for Inference

Column Dimension: Assuming that the entries of a data column are conditionally independent given the column dimension, we obtain the posterior distribution of column dimension t by marginalizing over row unit and label variables

\mathbf{u} and \mathbf{z} as follows:

$$\begin{aligned}
p(t = k|\mathbf{x}) &\propto p(t = k, \mathbf{x}), \\
&= p(t = k) \prod_{i=1}^N p(x_i|t = k), \\
&= p(t = k) \prod_{i=1}^N \left[\sum_{u_i=1}^{L_k} \sum_{z_i \in \{u_i, m, a\}} p(x_i, z_i, u_i|t = k) \right], \\
&= p(t = k) \prod_{i=1}^N \left[\sum_{u_i=1}^{L_k} p(u_i|t = k) \left(w_{u_i}^k p(x_i|z_i = u_i) \right. \right. \\
&\quad \left. \left. + w_{u_i}^m p(x_i|z_i = m) + w_{u_i}^a p(x_i|z_i = a) \right) \right]. \quad (1)
\end{aligned}$$

Eq. 1 can be used to estimate the column dimension t , since the one with maximum posterior probability is the most likely dimension corresponding to the column \mathbf{x} . As per the equation, the model estimates the column dimension by considering all the data rows, i.e. having missing or anomalous unit symbols does not confuse the dimension inference. Note that such entries would have similar likelihoods for each column dimension, which allows the model to choose the dominant dimension for regular entries.

Row Label: Following the inference of column dimension, the posterior probabilities of each row label z_i given $t = k$ and x_i is obtained by marginalizing the latent unit variable u_i as follows:

$$p(z_i = j|t = k, x_i) = \frac{p(z_i = j, x_i|t = k)}{\sum_{\ell \in \{u_i, m, a\}} p(z_i = \ell, x_i|t = k)},$$

where $p(z_i = j, x_i|t = k)$ can be calculated as follows:

$$\begin{aligned}
p(z_i = j, x_i|t = k) &= \sum_{l=1}^{L_k} p(u_i = l, z_i = j, x_i|t = k), \\
&= \sum_{l=1}^{L_k} p(u_i = l|t = k) w_l^j p(x_i|z_i = j). \quad (2)
\end{aligned}$$

Row Unit: Given $t = k$, $z_i = j$ and x_i , the posterior distribution of row unit u_i is obtained as:

$$p(u_i = l|t = k, z_i = j, x_i) = \frac{p(u_i = l, x_i|t = k, z_i = j)}{\sum_{u_i=1}^{L_k} p(u_i = l, x_i|t = k, z_i = j)},$$

where $p(u_i = l, x_i|t = k, z_i = j)$ can be calculated as follows:

$$\begin{aligned}
&= p(u_i = l|t = k) p(x_i|t = k, u_i = l, z_i = j), \\
&= p(u_i = l|t = k) p(x_i|z_i = j).
\end{aligned}$$

Column Unit: Following the column dimension inference, we set the column unit l^* as follows:

$$l^* = \operatorname{argmax}_l \sum_{i=1}^N p(u_i = l | t = k, x_i), \quad (3)$$

where $l \in \{1, \dots, L_t\}$ denotes a possible unit for dimension t .

$$\begin{aligned} p(u_i = l | t = k, x_i) &\propto p(u_i = l | t = k) p(x_i | u_i = l, t = k), \\ &= p(u_i = l | t = k) \left[\sum_{\ell \in \{u_i, m, a\}} \left(p(x_i, z_i = \ell | u_i = l, t = k) \right) \right], \\ &= p(u_i = l | t = k) \left[\sum_{\ell \in \{u_i, m, a\}} \left(p(z_i = \ell | u_i = l) p(x_i | z_i = \ell) \right) \right]. \end{aligned}$$

Correcting Anomalous Unit Symbols: Row label inference annotates each data entry either as a unit, missing or anomalous. We assume that the units of anomalous entries are encoded by anomalous unit symbols (e.g., `ltrs` for `litres`) and can be identified by mapping anomalous unit symbols (e.g., `ltrs`) to known unit symbols (e.g., `lt`) based on the edit-distance [12]. The edit-distance measures the minimum number of operations (addition, deletion or substitution) that needs to be done to transform a string to another, and can handle misspellings and non-standard abbreviations. Note that we restrict the set of unit symbols to be compared with according to the column dimension, i.e., we compare `ltrs` with known unit symbols that encode units of volume when the column dimension is inferred as volume.

Canonicalizing Units: Following the inference of the row units and the column unit, we are now interested in representing each row with the same unit by scaling its numerical value (e.g., converting the data entry `1 m` to `100 cm` when the column and row units are respectively centimetres and metres). Currently, we convert units via an existing tool named Pint [8]. See Sec. 3 for a detailed discussion of Pint.

D Additional Experimental Results

D.1 Dimension Inference

PUC correctly predicts the dimensions of all data columns, except one data column for which all the competitor methods also fail. These failures result from the inability to parse the data values such as `5'10"` where `'` and `"` denote respectively feet and inches. We could improve our regular expression to parse such formats, which we have not done in order not to optimise on test datasets. On the other hand, the leading competitor methods are S-NER and Quantum,

which achieve the same overall accuracy. They both fail to identify the column dimensions of seven data columns. To compare our method with these baselines, we present their normalised confusion matrices in Fig. 2, normalised so that a column sums to 1.




Fig. 2: Normalized confusion matrices for (a) Quantum, (b) S-NER and (c) PUC plotted as Hinton diagrams, where the area of a square is proportional to the magnitude of the entry.

Fig. 2(b) shows that S-NER tends to infer the column dimension as length. These failures can be explained by the differences in the number of characters of unit symbols. S-NER performs better on longer unit symbols, which may not be surprising as it uses n-grams for feature extraction. For example, the confusions between volume and length occur on the HES dataset, where the dimensions of units symbols for cubic foot (e.g., `cuft`, `cu.ft`) are correctly predicted as volume. Nevertheless, the dimensions of unit symbols for litre (e.g., `l`, `L`) are predicted as length instead of volume, which result in incorrect predictions since they are more frequent in the data. Character-level features (e.g., n-grams) could be useful to handle variations in the data such as misspellings; however, they may lead to limited performance on short unit symbols. This result indicates the advantage of incorporating knowledge about unit symbols directly into the model, as in PUC.

As we discuss above, Quantum cannot parse measurements such as `5'10"` and fails to identify the dimension of the corresponding data column. The confusion between length and data storage occurs in a data column where kilobyte and megabyte are respectively encoded by `K` and `M`. Here, Quantum mislabels `M` as metre and does not generate a prediction for `K`. A detailed inspection of the remaining five columns shows that numeric values are missing in the entries, which cannot be handled by Quantum.

To determine whether the column dimension predictions of PUC and the leading competitor methods (namely S-NER and Quantum) are significantly different, we apply a variation of the McNemar’s test as the number of sam-

ples is low (see e.g., [4]). This test assumes that the two methods should have the same error rate under the null hypothesis. We compute the exact p-value $2 \sum_{i=n_{10}}^{n_{01}} \binom{n}{i} 0.5^i (1-0.5)^{n-i}$ where $n = n_{01} + n_{10}$ with n_{01} and n_{10} which respectively denote the number of columns misclassified by only a competitor method, and by only PUC. The test to compare PUC and S-NER results in a p-value of 0.03 since n_{01} and n_{10} are respectively equal to 6 and 0. These results reject the null hypothesis that the means are equal and confirm that they are significantly different at the 0.05 level. We obtain the same result from the test between PUC and Quantum, as n_{01} and n_{10} are the same as S-NER, which confirms that they are significantly different at the 0.05 level.

Lastly, we have evaluated the runtime of each method per data column. Figure 3 shows that S-NER is the slowest method, whereas PUC is the fastest on average. The leading competitor method is Pint; however, the variation in its runtime is higher than PUC. Note that we have not explicitly optimised our method for speed, which may improve further its scalability.




Fig. 3: Runtime violin plots denote the time in seconds taken to infer dimensions per column. The dot, box, and whiskers respectively denote the median, interquartile range, and 95% confidence interval.

D.2 Unit Identification

PUC outperforms the competitor methods by a large margin on 4 datasets (143...23, HES, MBA and Zomato) out of the 15. Note that we exclude the Taser dataset from the evaluations since none of the methods could parse its values such as 5'10" where ' and " denote respectively feet and inches. On the remaining 11 datasets, there is at least one competitor method competitive with ours.

The performance gap between PUC and the competitor methods reflects the importance of mapping anomalous unit symbols to known symbols through string-similarity. For example, on the HES dataset, the competitor methods could accurately identify only a few unit symbols, whereas our method could successfully identify almost all of the unit symbols. Out of 14 unique unit symbols, CCUT identified `L` and `cuft`, and Pint identified `L`, `l` and `litres`. In addition to these three unit symbols, Quantulum identified `Litres`. Surprisingly, GQ, which is one of the state-of-the-art methods in identifying units in text documents, could not identify any of these unit symbols. PUC, on the other hand, could identify 12 unit symbols correctly, with only two unidentified unit symbols (`Cu` and `cf`).

We observe that Quantulum performs better than PUC on the Arabica dataset. This result is mainly due to the Altitude column where metre is encoded by `M`, which is a known symbol for mile. Consequently, our method predicts the units of such entries as mile rather than metre. We could avoid this confusion by making row units dependent, so that the presence of unit symbols (e.g., `m`, `metres`) in the other data entries that encode the same unit (e.g., `metre`) is treated as an indicator of `M` being a symbol for metre rather than mile. Here, we do not adapt our model accordingly so that it is not optimised on test datasets.

On the Zomato dataset, PUC performs better than the competitor methods. However, it still cannot identify 4 symbols out of 10 currency symbols and maps them to incorrect units based on string similarities. For example, PUC identifies `QR` (a symbol for Qatari Riyal) as an anomalous unit symbol. As `QR` is an out-of-dictionary unit symbol, it is mapped by PUC to the known `BRL`, which is a symbol for the South African Rand. Although this mapping mechanism can improve performance in certain cases (e.g., the HES dataset), it can sometimes fail.

To determine whether the performances of PUC and Pint are significantly different, we have applied a paired t-test on the differences of the accuracies, i.e., $\text{Accuracy}(\text{PUC}) - \text{Accuracy}(\text{competitor method})$. We have calculated the t-statistic of 4.85 and the p-value of 0.0002 for CCUT, the t-statistic of 8.61 and the p-value of 0.000001 for GQ, the t-statistic of 2.74 and the p-value of 0.01 for Pint and the t-statistic of 2.55 and the p-value of 0.02 for Quantulum. These results reject the null hypothesis that the means are equal and confirm that they are significantly different at the 0.05 level.

Lastly, Table 2 indicates that Pint and Quantulum perform best among the competitor methods. We have compared their performances through a paired t-test, i.e., $\text{Accuracy}(\text{Quantulum}) - \text{Accuracy}(\text{Pint})$, which suggests that they are not significantly different at the 0.05 level (the t-statistics of 0.18 and the p-value of 0.86). This result suggests that neither method is comprehensive enough to handle various unit symbols observed in different datasets.