

LaTeX-Numeric: Language-agnostic Text attribute eXtraction for E-commerce Numeric Attributes

Kartik Mehta
India Machine Learning
Amazon
kartim@amazon.com

Ioana Oprea
Retail Business Services
Amazon
ioanao@amazon.com

Nikhil Rasiwasia
India Machine Learning
Amazon
rasiwasi@amazon.com

Abstract

In this paper, we present LaTeX-Numeric - a high-precision fully-automated scalable framework for extracting E-commerce numeric attributes from product text like product description. Most of the past work on attribute extraction is not scalable as they rely on manually curated training data, either with or without the use of active learning. We rely on distant supervision for training data generation, removing dependency on manual labels. One issue with distant supervision is that it leads to incomplete training annotation due to missing attribute values while matching. We propose a multi-task learning architecture to deal with missing labels in the training data, leading to F1 improvement of 9.2% for numeric attributes over single-task architecture. While multi-task architecture benefits both numeric and non-numeric attributes, we present automated techniques to further improve the numeric attributes extraction models. Numeric attributes require a list of units (or aliases) for better matching with distant supervision. We propose an automated algorithm for alias creation using product text and attribute values, leading to a 20.2% F1 improvement. Extensive experiments on real world dataset for 20 numeric attributes across 5 product categories and 3 English marketplaces show that LaTeX-numeric achieves a high F1-score, without any manual intervention, making it suitable for practical applications. Finally, we show that the improvements are language-agnostic and LaTeX-Numeric achieves 13.9% F1 improvement for 3 Romance languages¹.

1 Introduction

E-commerce websites often sell billions of products. These websites provide information in form of product images, product text (such as title and product description) and structured information,

henceforth, termed as product attributes². These attributes often act as a concise summary of product information and are useful in product discovery, comparison and purchase decisions. They are usually provided by selling partners at the time of product listing and can be missing or invalid, even though they might be present in product text sources. Extracting attribute values from these product text sources can be used to populate the missing attribute values and is the focus of this work.

Attribute Extraction from free form text can be posed as a Named Entity Recognition (NER) problem (Zheng et al., 2018). Recently, deep learning models (Lample et al., 2016; Ma and Hovy, 2016; Huang et al., 2015) have shown remarkable performance on NER tasks, eliminating the need of manually curated features. However, these approaches still require large amount of labelled data. While active learning can be used to efficiently curate training data (Zheng et al., 2018), however gathering data for hundreds of product categories and attributes is a resource extensive task. One solution is to use distant supervision to create training data. Distant supervision has been extensively used to curate training set without manual effort for relation extraction (Mintz et al., 2009). In context of attribute extraction for E-commerce, we can curate training data by using attribute values and matching them with tokens in product text. However, if attribute values are missing, distant supervision leads to missing annotations, a phenomenon not studied in literature.

In this work, we present an automated framework for building high-precision attribute extraction models for numeric attributes using distant supervision. Multiple works in literature (Madaan et al., 2016; Ibrahim et al., 2016) argue that distant

¹<https://www.britannica.com/topic/Romance-languages>

²E.g. RAM, weight and front_camera are some of the product attributes for mobile phone. We use the terminologies 'product attributes' and 'attributes' interchangeably in this paper.

supervision for numeric attributes poses unique challenges and have given separate treatment to numeric attributes. Highlighted below are some interesting challenges that distant supervision poses for numeric attribute extraction models:

Partial Annotations: Distant supervision leads to incorrect annotations when attribute is present in the text field but structured attribute value is missing.

Diverse surface forms: There are multiple ways that attributes are mentioned in product text (e.g. resolution of ‘2’ can be mentioned as ‘2 mp’, ‘2 mpix’ or ‘2 megapixels’). We term these different surface forms as alias.

Confusing attributes: Many attributes have common units and may have confusing mention in the text (e.g. ‘16 GB memory’ refers to RAM while ‘128 GB memory storage’ refers to ‘Hard Disk’)

Use of different units: Seller may use diverse units for numeric attributes (e.g. ‘1.5 kg’ as attribute value and ‘3.3 pounds’ in product text).

Addressing these challenges in automated manner is the primary focus of this work. Our paper has the following contributions: (1) We propose a multi-task architecture to deal with partial annotations introduced due to missing attributes. This multi-task architecture leads to F1 improvement of 9.2% for numeric and 7.4% for non-numeric attributes over single task architecture. (2) We propose a fully automated algorithm for alias creation using product text and attribute values. These alias improve the quality of training annotation in distant supervision, leading to models with 20.2% F1 improvement for numeric attributes. We demonstrate the effectiveness of our proposed approach using a real-world dataset of 20 numeric attributes across 5 categories and 3 English marketplaces. Models trained using our proposed framework achieve a high F1-score without any manual intervention, making them suitable for practical applications. We show that our proposed approach is language agnostic. Experiments of using our framework on 3 Romance languages show 13.9% F1 improvement. To the best of our knowledge, this is first successful attempt at building automated attribute extraction for numeric attributes at E-commerce scale. Rest of the paper is organized as follows. We describe our proposed framework and its components in Section 3. We describe the ‘Multi Task’ architecture in Section 3.1 and ‘automated alias creation’ component in Section 3.2. We describe datasets,

experimental setup and results in Section 4. Lastly, we summarize our work in Section 5.

2 Related Work

2.1 Attribute Extraction for E-commerce

Early works on information extraction focused on extracting facts from generic web pages (Oren et al., 2005; Yates et al., 2007; Etzioni et al., 2008). With rise of E-commerce, multiple works focused on extracting attributes from product pages. Ghani et al. (2006) proposed use of supervised learning to extract attributes from E-commerce product descriptions. Putthividhya and Hu (2011) formulated attribute extraction from short titles as NER problem, using multiple base classifiers and a CRF layer. The training data was created by matching entries from a seed dictionary. More (2016) proposed use of distant supervision for attribute extraction. They used token-wise string matching (henceforth referred as exact match) based on attribute values to annotate title tokens and train an NER model with manually defined features. They used manual intervention to improve the training annotations e.g. dealing with spelling mistakes and different surface forms of brand. Majumder et al. (2018) extended this work with use of recurrent neural networks, excluding use of manually defined features. Zheng et al. (2018) proposed OpenTag using bidirectional LSTM, Conditional Random Fields (CRF) and attention mechanism. But the training data for OpenTag is manually created with use of active learning, making it challenging to use at E-commerce scale.

2.2 Distant supervision of numeric attributes

Getting manual training data has always been a resource intensive and expensive task and distant supervision has been explored as an alternative. Distant supervision for numeric attributes has been used for relation extraction (Hoffmann et al., 2010; Madaan et al., 2016), question answering (Davidov and Rappoport, 2010), entity linking (Ibrahim et al., 2016). Madaan et al. (2016) argued that distant supervision for numerical attributes presents peculiar challenges not found for non-numeric attributes, such as high noise due to matching out of context, low recall due to different rounding level, and importance of units. Ibrahim et al. (2016) constructed a KB from freebase.com, keeping a list of units and conversion rules for numeric quantities. While these works have established the im-

Attribute	Value
CPU speed	3.5 Ghz
RAM	16 GB
Flash Memory	128 GB
Hard Disk size	1 TB
Screen Size	15.6 inch
Battery Life	?

Intel Core i7-6700HQ **3.5 GHz**
 Processor with NVIDIA GeForce GTX 1060 with 6GB GDDR5
 Graphics Memory ## **16GB** DDR4 included; 32GB maximum ##
128GB SSD (Boot) + **1TB** 7200Rpm
 ## **15.6 inch** FHD (1920 x 1080) IPS
 Anti-Glare 300-nits Display [FHD]
 ## Windows 10 home operating system with Microsoft signature image; silver

Figure 1: Illustration of E-commerce attribute extraction problem.




Figure 2: LaTeX-Numeric framework for extraction of E-commerce numeric attributes.

importance of units for distant supervision of numeric attributes, but the list of units is manually curated.

2.3 NER with partial annotation

Distant supervision may lead to noisy training data due to partial annotations. Tsuboi et al. (2008) argued that partial annotations may happen due to ambiguous annotation and proposed CRF-PA to alleviate the issue of partial annotations. Yang et al. (2018) studied partial annotations introduced due to incomplete dictionary and extended CRF-PA to NER models. Jie et al. (2019) proposed learning the probability distribution of all possible label sequences compatible with given incomplete annotation, and using this probability to clean the training annotations. For E-commerce attribute extraction, partial annotations may happen due to missing attribute value. Our paper is the first work to establish this phenomenon for attribute extraction and provide a systematic way to alleviate this problem. We compare our proposed approach with Jie et al. (2019) in Section 4.2.

3 LaTeX-Numeric Framework

We pose the attribute extraction problem (refer Figure 1) as a NER problem, where product attributes are treated as named entities. Formally, we are given a text X with a particular tokenization (x_1, x_2, \dots, x_m) and a set of attributes $A: (\alpha_1, \alpha_2, \dots, \alpha_n)$. The task is to extract $v_i = \alpha_k$ for $i \in [1, m]$ where $k \in [0, n]$ and α_0 represents ‘Other’ entity.

Figure 2 gives an overview of our proposed LaTeX-Numeric framework. We are given a list of pre-defined numeric attributes, dump of prod-

ucts consisting of product text and existing attribute values. The attribute values are decimals (e.g. 16) and have an underlying unit (e.g. GB). We term this underlying unit as the canonical unit. For creating distant supervision-based training annotations, we use these canonical units and combine them with attribute values for matching with product text. This serves as the ‘canonical aliasing’ baseline for our comparisons. We use BIO tagging scheme for our experiments as it is a popular format. For training, we use the recently proposed BiLSTM-CNN-CRF model (Ma and Hovy, 2016). This model consists of CNN architecture to encode character information, LSTM-based encoder to model contextual information of each token and a CRF based tag decoder, which exploits the labels of neighboring tokens for improved classification. Unlike Open-Tag (Zheng et al., 2018), we don’t use attention as we didn’t observe any improvements with use of attention in our initial experiments.

Creating training annotations using distant supervision may lead to partial annotations due to missing attribute values. In Section 3.1, we describe our proposed multi-task learning architecture to deal with such partial annotations. Additionally, we have observed that sellers use multiple surface forms to mention attributes in product text (e.g. ‘3mp’, ‘3mpix’, ‘3 megapixels’ for resolution) and hence, distant supervision with just canonical units (e.g. ‘mp’) may lead to suboptimal training annotations. Curating a list of these diverse surface forms will help improve the quality of training annotations. We describe an automated approach for curating such diverse units and improving training annotations of numeric attributes in Section 3.2.

3.1 Multi Attribute Joint Extraction

To jointly extract multiple attributes, the tagging strategy can be modified to consider an output label with tags for all attributes. With ‘BIO’ tagging, each attribute has its own (‘B’ and ‘I’) tags with ‘O’ common for all attributes, leading to total $2K + 1$ tags for K attributes. Based on this modified tagging, a single NER model can be trained for multi-attributes extraction. We term this setting of training ‘Multi Attribute Single Task’ model as MAST-NER (refer Figure 3). MAST-NER is the commonly used strategy for attributes extraction (Zheng et al., 2018; Sawant et al., 2017; Shen et al., 2017; Joshi et al., 2015).

Under distant supervision, attribute value is used to find matches in product text tokens. However,

if the attribute value is missing, no match will be found even when attribute value is mentioned in text and hence, the corresponding tokens are incorrectly tagged as ‘O’. We term this partial annotation due to missing attribute as Missing-PA problem. Table 1 shows an illustration of this problem. Missing-PA is generic to distant supervision for multi-attributes and exists for non-numeric attributes as well. To the best of our knowledge, this problem has not received attention in literature.

	Display	RAM	Weight	BatteryLife
Attribute Value	12.3	16	missing	10
Canonical Unit	inches	gb	kg	hours

The high performance Chromebook. Features 7th Gen Intel Core i7 processor, **16 GB** RAM and 512 GB for storage. The long lasting battery provides up to **10 hours** of use and its fast charging so you can get 2 hours of use in 15 minutes ## Pixelbook’s super thin and lightweight design measures 10.3 mm and weighs **1.2 kg** Features a **12.3 inches** 360 touchscreen display

Table 1: Illustration of Missing-PA for distant supervision. 1.2kg will be incorrectly tagged as ‘O’ as value for weight attribute is missing.

To alleviate Missing-PA problem, one can train separate models for each attribute, by excluding samples where the corresponding attribute value is missing. However, such an approach requires training and managing a large number of models and separate computation for each attribute at evaluation time. Due to these practical challenges, this strategy is not suitable for practical applications. Another way to alleviate missing-PA is to use MAST-NER setting, and to exclude all samples where atleast one attribute has missing value. However, this approach may significantly reduce the size of training data as some attributes may have high missing rate, leading to a suboptimal model. To alleviate this problem, we propose a multi-task learning architecture with separate output layers for each attribute as separate tasks. We term this architecture of training ‘Multi Attribute Multi Task’ model as MAMT-NER (Refer Figure 3). MAMT-NER consists of shared character encoder, word encoder and BiLSTM layers. For each training sample, loss is deactivated (using masking) for tasks where corresponding attribute value is missing and activated only for remaining tasks where corresponding attribute values are non-missing. Loss for all activated tasks are weighted uniformly and weights of those tasks (including shared weights)

are updated for the given sample. Note that the proposed MAMT-NER architecture is generic and can be used for non-numeric attributes as well as any underlying NER architecture, including recently proposed BERT (Devlin et al., 2019).

3.2 Automated Alias Creation

As argued earlier, the canonical unit is often not sufficient to capture diverse surface forms that sellers use to mention attributes in product text. E.g. ‘13 inch’, ‘13 inches’, ‘13 in’, are multiple ways to mention display_size. One can analyze the mention of attribute values in product text and leverage that to create a list of commonly used surface forms. While, such an algorithm will detect common surface forms, it will miss out on units which require a multiplicative factor (e.g. ‘pounds’, ‘lbs’ and ‘ounces’ for weight where attribute values are in ‘kg’). To detect such units, we can analyze all numeric mentions in product texts (in isolation from attribute value) and filter out noisy candidates by using similarity with canonical units in embedding space. Additionally, we have observed that some numeric attributes have units which are specific to those attributes (e.g. ‘mah’ for battery_power and ‘hertz’ for refresh_rate). One can detect such attributes and use this information while creating training annotations using distant supervision. Based on these learnings, we propose an approach for generating a more exhaustive list of aliases, in an automated fashion (Figure 4).

3.2.1 Creation of alias_dw

We create attribute-specific alias_dw using product text and attribute values. We use a regex matching function, \mathcal{M}^3 , to find candidate alias which are preceded by numeric attribute value in text. Though this matching function may lead to instances of collision (e.g. 5 for RAM attribute may match with 5 ghz in text), we ignore cases where more than one match is found in text, to prevent impact of collisions. Alias_dw leads to common surface form of attribute units (e.g. ‘in’, ‘inches’, ‘inch’ for display_size and ‘gb’, ‘gigabyte’ for hard_disk).

3.2.2 Creation of alias_bp

We create a single alias_bp (common across attributes) using product text data. We use a regex function \mathcal{F}^4 to find candidate alias, which are to-

³ $\mathcal{M} = re.findall(" " + <value> + r"[-]*[a-zA-Z] + ", <text>)$, where <value> is attribute value (e.g. 8 for RAM) and <text> is product text.

⁴ $\mathcal{F} = re.findall(" " + r"[\d\.]*[\d] [-]*[a-zA-Z] + ", <text>)$, where <text> is product text.

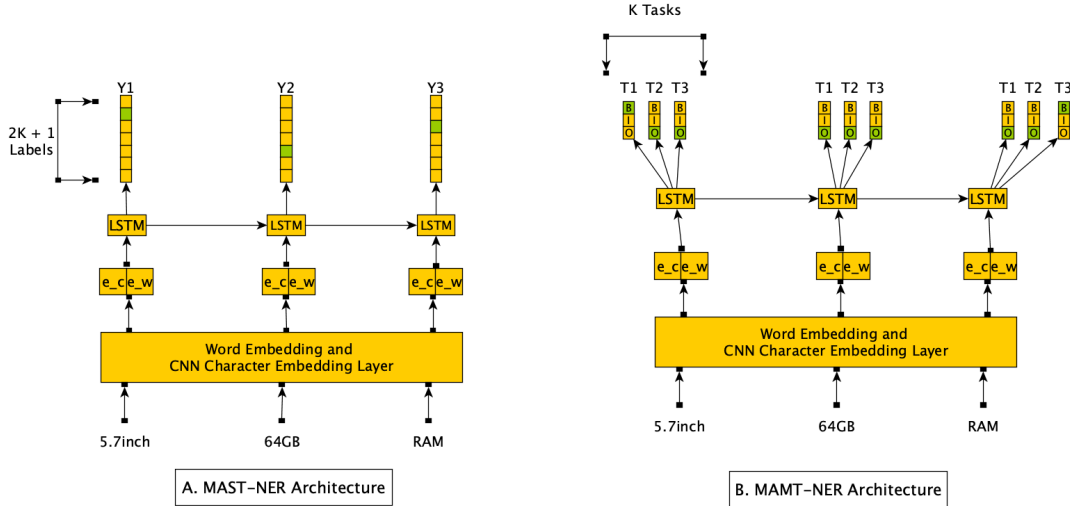


Figure 3: Figure showing different architectures for Multi Attributes Extraction models. We assume BIO-tagging of attributes with only 3 tags possible - B, I and O. For MAST-NER, we have two possible tags for each attribute and one others tag. For MAMT-NER, each attribute extraction is considered a separate task with weights shared for character embeddings, word embeddings and BiLSTM layer

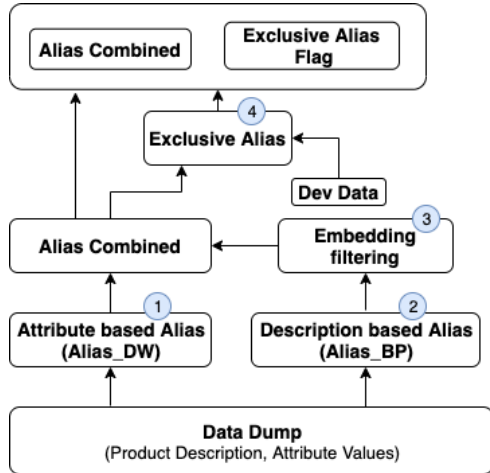


Figure 4: Flow-diagram for Automated alias creation.

kens preceded by any numeric mention in product text. Alias_bp may contain noisy candidates which are not units for any attribute. We use word embeddings to match alias_bp candidates with attributes and exclude noisy candidates.

3.2.3 Embedding based filtering

To remove noisy candidates and match alias_bp candidates to attributes, we leverage canonical units and Glove embeddings. For each attribute, we calculate similarity of each alias_bp candidate with its canonical unit in embedding space and keep only those candidates where similarity is greater than a pre-determined threshold. Thus, we obtain alias_bp_filter, which is attribute specific. E.g. we filter out ‘inches’ and select ‘pounds’ and ‘lbs’ for weight attribute having ‘kg’ as canonical unit.

Alias_dw and alias_bp_filter complement each other. Alias_bp_filter misses out on units which

have low similarity using embeddings (e.g. ‘in’ for display_size as ‘in’ has a low similarity with ‘inch’). Alternately, alias_dw misses out on cases where the unit mentioned in product text may require a multiplicative factor (e.g. alias_dw for item_weight misses out on ‘pounds’ and ‘lbs’). We concatenate alias_dw and alias_bp_filter to obtain alias_combined for each attribute (shown for four attributes in Table 2).

With small manual effort, one can get the multiplicative factor for converting values in canonical units to units in alias_combined and vice versa, which can further improve training annotations. As focus of current work is to build a fully automated attribute extraction system, we leave this as future work to be explored.

Category	Attribute	Alias_combined
Laptop	Hard_Disk	[gb, mb, gigabyte, tb, x]
Laptop	Display	[inches, inch, mm, cm, ft, centimeter, feet, in]
Tablet	Weight	[kilograms, kgs, kg, grams, lbs, pounds, ounces, g]
TV	Refresh_Rate	[hertz, hz]

Table 2: Alias values shown for few attributes

3.2.4 Exclusive Alias Flag

We use a small manually labelled dev set (created for hyper-parameters tuning) to create a flag indicating which attributes have exclusive alias. We evaluate precision of extracting any mention of alias for a given attribute and if this precision is above a threshold, we consider that attribute to have

exclusive alias. For attributes having an exclusive alias, we use regex-based matching for training annotations, tagging any numeric value followed by the corresponding unit, irrespective of the attribute value.

We refer our proposed approach of using ‘alias_combined’ and ‘exclusive alias flag’ for creating training data of numeric attributes as ‘auto-aliasing’ henceforth. We discuss experiments of using ‘auto-aliasing’ as compared to other distant supervision techniques for numeric attributes in Section 4.1.

4 Experimental Setup and Results

We picked five product categories and their 20 numeric attributes for three English marketplaces (IN, US and UK). We extracted product data (product description and attribute values) for these categories and split this data into two parts (80% train and 20% test). The train part is used for automated alias creation and creation of training annotations with distant supervision. From the test part, we randomly picked products for each category and label the mention of category-specific numeric attributes in text. Out of the total labelled attribute-product pairs, we observed mention of 6.9K attributes in product text. We term training data for English as ‘Train-EN’ and audited test dataset as ‘Test-EN’ (details in Table 3). To evaluate applicability of our proposed LaTeX-Numeric framework for non-English languages, we did a similar analysis with one product category for three Romance languages of French (FR), Spanish (SP) and Italian (IT). We term this training data as ‘Train-Romance’ and audited test dataset as ‘Test-Romance’. Similar to (Zheng et al., 2018), we use F1-score for evaluation. Predictions are given full credit if correct value is extracted, but extracting more values than actual is considered incorrect (e.g. for a mobile phone with ‘4 gb’ RAM, extracting either ‘4’ or ‘4 gb’ is considered correct prediction, but extracting two values of ‘4 gb’ and ‘16 gb’ is considered incorrect prediction).

4.1 Evaluation of Matching Techniques

In this section, we study improvements with our proposed alias creation. For comparison, we use two baselines of creating training annotation a) using lexical matching of numeric attribute value and product text (‘exact match’), and b) matching based on canonical units (‘canonical aliasing’). For each strategy, we use CNN-BiLSTM-CRF with MAST-

Product Category	IN	US	UK
A (6)	14K (854)	241K (966)	70K (468)
B (4)	11K (396)	57K (493)	171K (320)
C (6)	43K (1027)	112K (623)	76K (559)
D (2)	3K (201)	14K (146)	32K (90)
E (2)	11K (246)	89K (266)	41K (244)
Total (20)	83K (2724)	514K (2494)	391K (1681)

Product Category	FR	IT	ES
A (6)	93K (727)	70K (508)	74K (559)

Table 3: Stats for training and test data. Number of training products are shown with unit ‘K’ (K=1000) and number of labelled attributes mention in test data is shown in adjacent parenthesis. Number of attributes is shown in parenthesis adjacent to each category.

Matching Technique	IN	US	UK	Avg
exact match	78.0	86.5	93.3	85.5
canonical aliasing	100.0	100.0	100.0	100.0
auto aliasing (our)	113.1	120.3	128.5	120.2

Table 4: Comparison of various matching techniques for training data generation using distant supervision (all numbers are relative to using canonical units).

NER architecture. Table 4 shows F1 score for using different matching techniques. ‘Canonical aliasing’ approach shows better F1 score than ‘exact match’, but it still suffers from low recall due to missing out on different surface forms mentioned in product text. With our proposed auto-aliasing, we address this limitation of ‘canonical aliasing’ and observe an average F1 improvement of 20.2%, establishing ‘auto-aliasing’ as best technique for distant supervision of numeric attributes. We use the training data created using ‘auto-aliasing’ for all subsequent experiments (unless otherwise specified).

4.2 Evaluation of MAMT Architecture

In this section, we perform a quantitative evaluation of our proposed MAMT architectures. Table 5 shows results of using MAST and MAMT architecture with CNN-BiLSTM-CRF. As compared to MAST-NER architecture, we observe 9.2% F1 improvement with our proposed MAMT architecture. Additionally, we observe that Jie et al. (2019)⁵ shows better F1 score than MAST due to higher recall. However, Jie et al. (2019) leads to drop in precision due to confusion between close attributes

⁵We use implementation of https://github.com/allanj/ner_incomplete_annotation. We show results only for IN as we get memory error when training for US and UK datasets which have larger training size.

Model + Architecture	IN	US	UK	Avg
BiLSTM-MAST	113.1	120.3	128.5	120.2
BiLSTM (Jie et al. (2019))	114.4	NA	NA	NA
BiLSTM-MAMT	124.4	131.4	139.0	131.2
BERT-MAST	117.7	122.9	128.8	122.8
BERT-MAMT	120.4	127.8	134.3	127.1

Table 5: Study of multi-task architecture for numeric attributes. BERT uses softmax as output layer, while, BiLSTM refers to CNN-BiLSTM model with crf as output layer. All numbers are relative to using canonical units in Table 4.

Model	Architecture	Precision	Recall	F1
BiLSTM	MAST	100.0	100.0	100.0
BiLSTM	MAMT	93.6	117.8	107.4

Table 6: Study of multi-task architectures for textual attributes (all numbers are relative). BiLSTM refers to CNN-BiLSTM model with crf as output layer.

(e.g. front-camera and back-camera for mobile.) Our proposed MAMT shows 8.7% better F1 score than Jie et al. (2019) for IN.

Further, we do comparison of MAST and MAMT architectures with BERT⁶ as underlying model and observed 3.5% F1 improvement with MAMT, demonstrating its applicability to multiple underlying models. To establish the effectiveness of MAMT architecture for non-numeric attributes, we curated a test dataset of 600 samples per attribute for 8 textual attributes across 4 product categories. As shown in Table 6, we observe 7.4% F1 improvement on this dataset with our proposed MAMT-NER architecture, showing its effectiveness on textual attributes as well.

4.3 Evaluation on non-English Languages

Architecture + Matching	FR	IT	ES	Avg
MAST + canonical-aliasing	100	100	100	100
MAST + auto-aliasing	103.0	107.1	108.4	106.0
MAMT + auto-aliasing	104.0	118.4	121.6	113.9

Table 7: Comparison of auto-aliasing and multi-task architecture on Romance languages (numbers are relative to using canonical-aliasing).

In this section, we study the applicability of LaTeX-Numeric for three Romance languages. We train

⁶We use implementation of <https://github.com/namisan/mt-dnn>, which uses bert-base and softmax as output layer.

a separate (Category-A) model for each Romance language replacing English word embeddings with language specific fastText (Lample et al., 2018) embeddings. Table 7 shows results on Test-Romance dataset. We observe 6.0% F1 improvement with our proposed auto-aliasing and additional 7.9% improvement with use of MAMT-NER architecture, showing effectiveness of our proposed approaches across languages.

5 Conclusion

In this paper, we described ‘LaTeX-Numeric’, a high-precision fully-automated framework for training attribute extraction models for E-commerce numeric attributes. We characterized the problem of Missing-PA that arises with distant supervision due to missing attribute values and proposed a multi-task learning architecture to alleviate the Missing-PA problem, leading to 9.2% F1 improvement for numeric attributes. We established the applicability of our proposed multi-task architecture for textual attributes and BERT as underlying model as well. Additionally, we proposed an automated algorithm for alias creation, to deal with variations of numeric attribute mentions, leading to models with 20.2% F1 improvement. Our evaluation on three Romance languages establishes that these improvements are applicable across non-English languages as well. Models trained using our proposed LaTeX-Numeric framework achieve high F1 score, making them suitable for practical applications.

References

- Dmitry Davidov and Ari Rappoport. 2010. Extraction and approximation of numerical attributes from the web. In *48th Annual Meeting of ACL*, pages 1308–1317.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Oren Etzioni, Michele Banko, Stephen Soderland, and Daniel S Weld. 2008. Open information extraction from the web. *Communications of the ACM*, 51(12):68–74.
- Rayid Ghani, Katharina Probst, Yan Liu, Marko Krema, and Andrew Fano. 2006. Text mining for product attribute extraction. *ACM SIGKDD Explorations Newsletter*, 8(1):41–48.

- Raphael Hoffmann, Congle Zhang, and Daniel S Weld. 2010. Learning 5000 relational extractors. In *48th Annual Meeting of the ACL*, pages 286–295.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv*.
- Yusra Ibrahim, Mirek Riedewald, and Gerhard Weikum. 2016. Making sense of entities and quantities in web tables. In *25th CIKM*, pages 1703–1712. ACM.
- Zhanming Jie, Pengjun Xie, Wei Lu, Ruixue Ding, and Linlin Li. 2019. Better modeling of incomplete annotations for named entity recognition. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 729–734, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mahesh Joshi, Ethan Hart, Mirko Vogel, and Jean-David Ruvini. 2015. Distributed word representations improve ner for e-commerce. In *1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 160–167.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the NAACL-HLT*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations*.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-task deep neural networks for natural language understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- Aman Madaan, Ashish Mittal, Ganesh Ramakrishnan, Sunita Sarawagi, et al. 2016. Numerical relation extraction with minimal supervision. In *Thirtieth AAAI*.
- Bodhisattwa Prasad Majumder, Aditya Subramanian, Abhinandan Krishnan, Shreyansh Gandhi, and Ajinkya More. 2018. Deep recurrent neural networks for product attribute extraction in ecommerce. *arXiv*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Joint Conference of the 47th Annual Meeting of the ACL and the 4th IJCNLP of the AFNLP*, pages 1003–1011.
- Ajinkya More. 2016. Attribute extraction from product titles in ecommerce. *CoRR*, abs/1608.04670.
- Etzioni Oren, Cafarella Michael, Downey Doug, Popescu Ana-Maria, Shaked Tal, Soderland Stephen, Weld Daniel S, and Yates Alexander. 2005. Unsupervised named-entity extraction from web: An experimental study. *Artificial intelligence*, 165(1):91–134.
- Duangmanee Pew Putthividhya and Junling Hu. 2011. Bootstrapped named entity recognition for product attribute extraction. In *EMNLP*, pages 1557–1567. ACL.
- Nils Reimers and Iryna Gurevych. 2017. Optimal hyperparameters for deep lstm-networks for sequence labeling tasks. *arXiv*.
- Uma Sawant, Vijay Gabale, and Anand Subramanian. 2017. E-fashion product discovery via deep text parsing. In *26th International Conference on WWW*, pages 837–838.
- Yanyao Shen, Hyokun Yun, Zachary Lipton, Yakov Kronrod, and Animashree Anandkumar. 2017. Deep active learning for named entity recognition. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 252–256, Vancouver, Canada. Association for Computational Linguistics.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *54th Annual Meeting of the ACL (Volume 2: Short Papers)*, pages 231–235.
- Yuta Tsuboi, Hisashi Kashima, Hiroki Oda, Shinsuke Mori, and Yuji Matsumoto. 2008. Training conditional random fields using incomplete annotations. In *22nd COLING*, pages 897–904. ACL.
- Yaosheng Yang, Wenliang Chen, Zhenghua Li, Zhengqiu He, and Min Zhang. 2018. Distantly supervised ner with partial annotation learning and reinforcement learning. In *27th COLING*.
- Alexander Yates, Michael Cafarella, Michele Banko, Oren Etzioni, Matthew Broadhead, and Stephen Soderland. 2007. Textrunner: open information extraction on the web. In *Human Language Technologies: The Annual Conference of the NAACL: Demonstrations*, pages 25–26. ACL.
- Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. In *24th ACM SIGKDD*, pages 1049–1058. ACM.