



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/192543/>

Version: Published Version

Article:

Zhang, Z. and Song, X. (2023) An exploratory study on utilising the web of linked data for product data mining. *SN Computer Science*, 4 (1). 15.

<https://doi.org/10.1007/s42979-022-01415-3>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



An Exploratory Study on Utilising the Web of Linked Data for Product Data Mining

Ziqi Zhang¹ · Xingyi Song²

Received: 11 September 2021 / Accepted: 16 September 2022
© The Author(s) 2022

Abstract

The Linked Open Data practice has led to a significant growth of structured data on the Web. While this has created an unprecedented opportunity for research in the field of Natural Language Processing, there is a lack of systematic studies on how such data can be used to support downstream NLP tasks. This work focuses on the e-commerce domain and explores how we can use such structured data to create language resources for product data mining tasks. To do so, we process billions of structured data points in the form of RDF n-quads, to create multi-million words of product-related corpora that are later used in three different ways for creating language resources: training word-embedding models, continued pre-training of BERT-like language models, and training machine translation models that are used as a proxy to generate product-related keywords. These language resources are then evaluated in three downstream tasks, product classification, linking, and fake review detection using an extensive set of benchmarks. Our results show word embeddings to be the most reliable and consistent method to improve the accuracy on all tasks (with up to 6.9% points in macro-average F1 on some datasets). Contrary to some earlier studies that suggest a rather simple but effective approach such as building domain-specific language models by pre-training using in-domain corpora, our work serves a lesson that adapting these methods to new domains may not be as easy as it seems. We further analyse our datasets and reflect on how our findings can inform future research and practice.

Keywords Linked data · Web of data · Schema.org · Natural language processing · Nlp · Data mining · Product mining

Introduction

Recent years have seen significant increase in the adoption of the Linked Open Data practice [1] by data publishers on the Web. LOD refers to the practice of describing structured

data using standard markup languages (e.g. RDFa¹) and universal vocabularies (e.g. schema.org) that allow defining properties and relations of data, and publishing and inter-linking such data on the open Web (hence ‘semantic markup data’). While early LOD datasets primarily took form of a graph database such as DBpedia,² an increasingly popular decentralised approach has been the embedding of semantic markup data within HTML pages. The Web Data Commons³ (WDC) project extracts such markup data from the

✉ Ziqi Zhang
ziqi.zhang@sheffield.ac.uk
Xingyi Song
x.song@sheffield.ac.uk

¹ <https://www.w3.org/TR/rdfa-primer/>.

² <https://wiki.dbpedia.org/>.

³ <http://webdatacommons.org>.

¹ Information School, University of Sheffield, Regent Court, 211 Portobello, Sheffield, South Yorkshire S1 4DP, UK

² Department of Computer Science, University of Sheffield, Regent Court, 211 Portobello, Sheffield, South Yorkshire S1 4DP, UK

CommonCrawl⁴ as RDF n-quads,⁵ and release them on an annual basis. As by October 2021, its statistics showed that over 60% of web pages, or over 50% of websites contained semantic markups, amounting to over 82 billion quads—more than doubled that from 2019.

Such semantic markup data allow the ‘machine reading’ of the Web. It has not only driven the development of new products and data integration services, but also created unprecedented opportunities for research in the areas of Natural Language Processing (NLP, e.g. [2, 3]). This is because the RDF n-quads extracted from such data are described by universal vocabularies that define concepts, their properties and relationships. One of the most popular vocabularies is schema.org, which currently contains nearly 800 concepts and 1400 properties, and is used by over 10 million websites.⁶ Studies have shown that such data can be used to train models for various NLP tasks, such as event extraction [2] and entity linking [4].

A particular domain that is witnessing the boom of semantic markup data is e-commerce, where online shops are increasingly embedding product markup data described using schema.org vocabularies into their web pages in order to improve content accessibility. Among the over 82 billion RDF n-quads mentioned above, nearly 17% are related to products, and are described by schema.org vocabularies. As examples, previous studies [4, 5] showed that among all product offers, 95% had an n-quad related to their names, 65% had one for their description, 35% had one for their brand, and less than 10% had one for their category. Such product markup data can be potentially useful as language resources to benefit various product data mining tasks, particularly in the current trend towards employing neural network models in developing large language models that prove to be effective for a wide range of NLP tasks [6–10]. However, we identify a gap in existing studies in this area. While there have been a limited number of sporadic studies [4, 5, 11] towards this direction in the product domain, the data sources used, the processes applied to use these data sources, and the findings have been inconsistent. This has made it difficult to properly compare or evaluate the usefulness of such data, or even answer the question ‘to what extent can we, and how can we, exploit this gigantic semantic markup data for product data mining tasks’.

To address this gap, in this work, we explore a series of questions in the context of the product domain, which is chosen for two reasons. First, it is one of the most promising domains where a ‘critical mass’ of such resources has been created. Second, it is a domain that continues to garner interest from both researchers and practitioners, as evidenced by

a series of workshops and shared tasks sponsored by industries [11, 12]. We study the following questions:

- How can we transform such semantic markup data to potentially useful language resources?
- How useful are these language resources for downstream NLP tasks in the product domain?

To answer these questions in this work, we aim to achieve the following objectives:

- To investigate and develop different methods for building language resources out of the semantic markup data;
- To evaluate the created language resources on a number of product-related NLP tasks using state-of-the-art (SoTA) benchmarks;
- To develop an understanding of if and how quality issues in the semantic markup data may impact on the language resource creation;
- To outline the implications of our study and future research directions that may help advance research and practice in this area.

Methodologically, we start with processing the semantic markup data to transform them into different types of language resources that can be used for downstream NLP tasks. For this, we review SoTA research in the area of product data mining and identify three ways: training word-embedding models (e.g. [11]), continued pre-training of BERT-like language models (e.g. [9]), and training machine translation models (e.g. [13]) that are used as a proxy to generate product-related keywords. Next, we apply these language resources to three product-related NLP tasks and evaluate them using SoTA methods on current benchmarks: product classification, product linking, and fake product review detection. The exact methods are dependent on the task and the language resources used, and therefore, will be explained in detail later.

The novelty and originality of this work lies in that, it is the first study that systematically studies and evaluates the use of semantic markup data for multiple product data mining tasks. Therefore, our findings serve as useful references for future research in this direction. First, we report that among the three methods, only word embeddings are the most consistent method to improve the accuracy on all three tasks. The BERT language models and the MT-based product keywords on the other hand, do not bring consistent improvements, even though there have been many studies that successfully developed in-domain BERT models following the simple principle of continued pre-training a generic BERT using large domain-specific corpora. Thus our results serve a lesson that this method may not be as easy as it seems. Second, we conduct a number of analyses of the data

⁴ <https://commoncrawl.org/>.

⁵ <https://www.w3.org/TR/n-quads/>.

⁶ <https://schema.org/>.

and show that the biased domain representation in the data and lack of vocabulary coverage may have been attributing factors. In particular, methods such as BERT language modelling and machine translation may be potentially more susceptible to such ‘data quality’ issues than word-embedding modelling. Finally, we discuss how these findings can inform future research and practice, and we contribute our data as public resources which can be obtained upon request.⁷

We organise the remainder of this paper as follows. The next section reviews related work. Then, the Sects. “**Product Classification**” to “**Fake product review detection**” present our exploration on each of the three tasks. Each section will introduce our methodology, experiments, and presents the results. The “**Further analysis**” section looks at potential quality issues of the data. In the “**Discussion**” section, we discuss our findings and their implications, and this is followed by a conclusion in the last section.

Related Work

While our work belongs to the general field of data mining [14–16], to avoid stretching our literature review too thinly, here we define our criteria of literature selection. More generally, semantic markup data is a type of LOD resources and there have been a large number of studies [17, 18] and organised events⁸ on the creation and consumption of LOD. However, previous studies have predominantly looked at LOD resources that are published as a graph database such as the DBpedia and Wikidata, while very few focussed on the LOD published as semantic markup data embedded within web pages. The fundamental difference between the two is quality, which underpins the approaches that one can take to use such resources. Most of the LOD graph databases are well-curated, documented and maintained. Semantic markup data however, can be highly heterogeneous, noisy, and unbalanced [19]. There are also a blend of studies that focussed on creating further LOD resources out of existing ones [20], compared to those that actually use LOD as language resources for downstream language processing tasks. Our literature review therefore has a specific focus on the following areas: (1) work that uses semantic markup data to create language resources, and (2) work on the three tasks we focus on, i.e. product classification, linking, and fake product review detection. While our work is also broadly related to the use of neural networks in creating domain-specific language models, such as the BioBERT [6], Clinical BERT [7], SciBERT[8], E-Bert[9], and SMedBERT[10], these studies did not use semantic markup data and therefore, we do not expand our literature review to this broad

area as that would significantly increase the scope of our discussion. However, in Section 8, we discuss our results with respect to the findings of these earlier studies.

Semantic Markup Data as Language Resources

Research on using semantic markup data for downstream language processing tasks has just taken off in the recent years and therefore, studies addressing the creation of language resources from such data are limited. Primpeli et al. [21] adopted an unsupervised approach to create a very large training dataset for product entity linking using semantic markup data extracted from the 2017 Common-Crawl corpus. The process started with extracting product offers that contain product identifiers annotated using the schema.org vocabulary. Then offers with the same identifiers are placed in the same cluster, followed by a cleaning process to eliminate potentially noisy clusters. The end clusters are considered to be product offers referring to the same product entity, and are used to train entity linking models. The work is further extended in later studies by [4, 22, 23], and it was shown that this automatically created training dataset has a high quality and can be used to effectively train product entity matchers at high accuracy. While these studies investigated ad-hoc usages of semantic markup data as training data for specific tasks, our work explores possibilities of utilising such data to create language resources that are usable by a wider range of tasks.

The same authors also used a product corpus to train a domain-specific word-embedding model in [22] using the fastText model. Specifically, they extracted the brand, name and description properties annotated by schema.org from the same corpus above, to create a text corpus that is used to train fastText embeddings. This domain-specific embedding model gained minor improvement over a generic fastText embedding model on some product linking tasks. In comparison, our work also explores using product related corpora to train word-embedding models. However, we study if this can generalise to other product data mining related tasks.

Work that uses semantic markup data to train embedding models can be traced back to [24], where authors used schema.org annotations (names and descriptions) of product entities to train entity embeddings using the paragraph2vec model [25]. This approach suffers similar limitations as the above, where the embeddings learned in such a way are ad-hoc and can only be used for entities found in the training process of the embedding models. Our study explores more generic ways of learning word embeddings.

In the 2020 Semantic Web Challenge on product data mining (MWP2020, [11]), a corpus of 1.9 billion words extracted from the descriptions of product entities annotated by schema.org vocabulary was used to train word-embedding models. Compared to generic word embedding models,

⁷ <https://bit.ly/2MGpBR2>.

⁸ <https://events.linkeddata.org/ldow2018/>, <https://ld4ie.github.io/>.

such models contributed to better results on the product classification task when used with a fastText baseline [12]. However, they were not used by any of the participating teams of the shared task. This study fills this gap by thoroughly evaluating them on several product data mining tasks.

Product Classification

Product classification is typically treated as an entity classification task. The process involves extracting product metadata for feature representation, then train a supervised algorithm that learns to assign category labels (i.e. classes) to product instances based on their features. Since most existing methods follow a similar process by mainly differ in terms of the metadata used, feature representation methods, and machine learning algorithms, below we summarise related work from these angles and highlight their similarities and differences instead of discussing each individual method in detail.

Metadata To classify products, features must be extracted from certain product metadata. Rich, structured metadata are often not available. Therefore, the majority of literature have only used product names, such as [26–29] and all of those participated in the 2018 Rakuten Data Challenge [12]. Several studies used both names and product descriptions [13, 30–35], while a few used other metadata such as model, brand, maker, etc., which need to be extracted from product specification web pages by an Information Extraction process [24, 36]. In addition, [24] also used product images. The work by [5, 37] used product categories allocated by the vendors and embedded as semantic markup data within the web pages. To differentiate these from the classification targets in such tasks, we refer to these as ‘site-specific product labels’ or ‘categories’. The authors noted that despite the highly heterogeneous nature of such site-specific labels across different websites, they are still very useful for supervised classification. In comparison, this work explores a ‘new’ type of metadata - product-related keywords generated by a machine translation model trained on the massive product corpora. Compared to product metadata already existing in a dataset and are comparatively better quality, such keywords may be very noisy. Our work will be first to explore if these keywords generated in such a way can be useful for product classification.

Feature representation Generally speaking, for text-based metadata, there are three types of feature representation. The first is based on Bag-of-Words (BoW) or N-gram models, where texts are represented based on the presence of vocabulary in the dataset using either 1-hot encoding or some weighting scheme such as TF-IDF [27, 29, 30, 36]. This often creates high-dimensional sparse feature vectors. The second uses pre-trained word embeddings or Language

Models (LM) to create a relatively low-dimensional, dense feature vector of the input text.

Certain techniques will need to be applied in order to compose embeddings for long text passages based on single words, such as [26, 32] that computed text embeddings based its composing words, and [38] (non-product domain) and [37] that joined word-embedding vectors to create a 2D tensor to represent the text. In the more recent work that uses pre-trained LMs such as BERT (e.g. in [39]), the construction of text passage embeddings is taken care of dynamically by passing the input text through the language models directly, which will take into account the context of words. The third applies a separate learning process to learn a continuous distributional representation of the text directly from the downstream training datasets [24, 31, 32]. Our work will make use of feature representation methods from the second and the third types. However, studies by [26, 32, 37, 39] used general purpose, pre-trained word embeddings while we study the effects of embeddings purposefully trained on product-related corpora. Studies by [24, 31, 32] created ad-hoc representations of product entities discovered in the training set, while such representations cannot be generalised to other data or tasks, our study explores more generic, data-agnostic methods for composing such representations.

Algorithms The large majority of work has used supervised machine learning methods. These include those that use traditional machine learning algorithms [5, 24, 27, 30–32], and those that apply DNN-based algorithms [12, 28, 36, 37] usually based on CNN or RNN. These include the majority of the participating systems in the 2018 Rakuten Data Challenge. Besides, [5] also explored unsupervised methods based on the similarity between the feature representations of a product and target classes. [29] studied product clustering, which does not label the resulting product groups. These represent unsupervised methods. Further, [13] studied the problem as a machine translation task, where the goal is to learn the mapping between a sequence of words from product names, to a sequence of product categories. MWPD2020 [11] showed a trend towards using the pre-trained LMs for classification, such as those based on the BERT model [40]. All participants but one at the product classification task at MWPD2020 used LM-based classification methods. The work by [39] for example, used an ensemble model combining 17 different variants of the BERT model to achieve the best result on this task.

This study does not focus on developing novel algorithms but instead, reuse existing ones such as the fastText baseline in [12] and the DNN structure in [37]. It may however, reveal which algorithms are more susceptible to the different language resources created by this study.

Product Linking

Product linking or matching is the task of determining if multiple product offers found from different websites (sometimes even from the same website) refer to the same, identical product entity. Product linking can be achieved by one of the three approaches: classification (e.g. [11]), where product offer pairs are created a priori and are classified to match or non-match; clustering, where a dataset of product offers are split into groups and members within the same group are considered to be about the same product; or retrieval (e.g. [34]) where the goal is to find the matching product entity from an existing database for a given product offer. In both classification and retrieval, often a ‘blocking’ process is applied to reduce the search space. All three approaches depend on the calculation of ‘similarities’ between product offers and this makes use of product metadata. A good literature review on product linking can be found in [24]. Here, we summarise them in terms of metadata, feature representation, and algorithms in a similar fashion as before.

Metadata Similar to product classification, typically product linking will make use of product names (e.g. [34, 41–46]) and descriptions (e.g. [24, 34, 47]). The difference however, is that the task also makes use of a diverse range of structured product attributes (e.g. [34, 44, 45, 48]), often defined as ‘key-value’ pairs such as those that can be extracted from product specifications (e.g. product ID, model, brand, manufacturer). Intuitively, offers that have the similar sets of key-value pairs are more likely to match. Since such structured key-value attributes are often unavailable, many studies focussed on how to extract them from the descriptions of an offer [41, 49], or from the specification table of the source web page [24]. A small number of studies [24, 50, 51] also used product images. Similar to product classification, we will explore the usefulness of product-related keywords generated by a machine translation model trained on the product semantic markup data. This has not been explored before.

Feature representation Again, similar to product classification, broadly speaking, transforming textual metadata into feature representations is typically based on BoW (e.g. [43, 44]), pre-trained word embeddings or language models (e.g. [23, 24, 34, 45, 46, 52]), or learning word embeddings on the spot from the downstream task datasets (e.g. [45]). However, depending on the types of metadata, different methods may be adopted and then combined [49]. For example, structured key-value attributes are often kept as-is and compared as a BoW, particularly if the values are short (e.g. product IDs). In [44], a concept of ‘q-gram’ was introduced to represent short texts (especially key-value pairs) as a set of character n-grams. Longer texts such as descriptions are better represented using word embeddings or through LMs. In this direction, similar sets of methods to

product classification are used. For image data, typical pixel-based image representation approach is widely used [24, 50, 51]. Same as product classification, in terms of novelty, our work focuses on evaluating the word embeddings purposefully trained on product-related corpora while many earlier studies used generic word embeddings. We also use more generic methods for composing feature representations for product linking while previous models tried to learn ‘ad-hoc’ representations.

Algorithms Since the prediction of linking/matching of product offers depends on a notion of ‘similarity’, some methods will have an ‘intermediary’ step that converts product metadata features to similarity features [34, 43, 48]. This is typically done by applying similarity metrics—usually based on string form, or word/character distribution—to the textual feature representations of two offers. Again, depending on the metadata, different similarity metrics may be applied [34, 43, 44, 49]. This ‘intermediary’ process creates a feature vector consisting of similarity scores computed by different measures, or using different features. The vector is then subject to another process to determine if the two offers should match. However, as mentioned before, some methods [45, 48] do not require such an intermediary step, as the similarity computation is embedded as part of the method that tackles the task in an ‘end-to-end’ fashion.

In terms of the method to the end-task, most studies are based on supervised binary classification, which aims to determine if a pair of offers should match or not. Following a similar pattern to product classification, the classification algorithms have evolved from the traditional [34, 41, 47], to DNN-based [34, 45], to LM-based [23, 33, 46, 52]. Depending on the classification algorithm, the input could be the similarity feature vector of a pair (e.g. [41]) computed by the intermediary step, or directly the feature vectors derived from the metadata of each offer [23, 45, 52]. In the study by [23] which extends their earlier work in [52], the authors proposed a multi-task learning neural network based on the BERT model, tailored for the product linking task. In addition to learning to predict if two pairs of product offers refer to the same entity (binary classification), the model at the same times learns to predict the shared product identifier by the two offers (multi-classification). Additionally, one can also make use of cutoff thresholds of similarity to determine match/non-match [42].

Clustering is used in a number of studies, such as [44] that clustered offers based on the ‘q-grams’ derived from their names and key-value pairs; and [53] where a ‘strength of ties’ style of clustering was applied to a ‘network’ of important words derived from a pair of product offers to determine if they form a cohesive ‘community’ and therefore, should match.

Methods that require offer pairs as input will often require a ‘blocking’ pre-process that aims to reduce the search space

for pairs, to create a minimal set of pairs for classification. Blocking strategies are varied and often lightweight, such as [49] that is based on matching manufacturers and categories, and [42] that is based on string prefix.

A unique direction of research in product linking looks into automated expansion of training data, either in terms of training instances, or metadata that can be used for feature extraction. For example, [42] enriched the name of product offers with tokens retrieved using a web search engine. [33] used product offer names to fetch similar entities from Wikidata, to create additional training instances.

Compared to state-of-the-art, we focus on the sub-task of supervised, binary classification of match/non-match, while ignoring the ‘blocking’ process. Our method will use state-of-the-art algorithms, as our research focus is on evaluating the impact of the language resources created from the semantic markup data on existing algorithms.

Fake Product Review Detection

Fake reviews, as per [54], generally refer to reviews created in an attempt to mislead consumers (either in a positive or negative way). They are also known as deceptive opinions, spam opinions, or spam reviews [55]. Fake online reviews in e-commerce significantly affect consumers, merchants, and market dynamics. In extreme cases, they led to financial loss for companies and legal cases [56]. While traditionally, fake reviewers are written by humans, with the advancement of Natural Language Generation technology, it has been shown that fake reviews automatically generated by programs are even more difficult for human annotators to detect [57]. There is an extensive amount of studies on automated fake review detection and for that reason, we refer readers to the surveys by [54] while below we present a brief overview of this field, highlighting the novelty of our work. Further, in addition to studies focussing on detecting the content, there are work (e.g. [58, 59]) that detect spammers (users) and spammer groups (network) which we do not cover here.

Detecting fake review is predominantly treated as a supervised, binary text classification task. Thus similar to product classification, it involves extracting features of the review text (metadata), representing them in a machine processable format (feature representation), and training a model that is able to generalise patterns using the features and apply the patterns to unseen data (algorithm). In terms of features (metadata), [57] broadly categorised them into ‘lexical’ and ‘non-lexical’. Lexical features are attributes derived from text, such as words, n-grams, punctuations and latent topics. Non-lexical features are metadata related to the reviews (e.g. ratings, stars) or their authors (ID, location, number of reviews generated). In terms of feature representation and algorithms, same patterns to that of product classification are noticed due to the two tasks been handled by text

classification approaches. Briefly, research has evolved from early methods that use manually engineered features in a 1-hot encoding (e.g. [60]) to pre-trained word embeddings (e.g. [61, 62]), and learning representations of the target dataset as part of the model, on the spot (e.g. [63]). The use of machine learning algorithms also evolved from the earlier classic algorithms such as SVM and logistic regression (e.g. [60]), to deep neural networks (e.g. [61, 62]), to using very large LMs such as BERT (e.g. [57]).

Compared to the previous studies, our work does not aim to introduce new features or algorithms. Instead, we explore the usefulness of the feature representations learned from massive product-related semantic markup data on the task of fake review detection. Earlier work such as [62, 64] used word embeddings pre-trained on general purpose corpora, and [61] trained domain-specific word embeddings using an Amazon product review corpus. In contrast, our work is the first that explores if a corpus of product details (instead of their reviews) can be used to learn word embeddings for this task. Compared to [57] who also used LMs, our work explores the effect of continued pre-training of LM using in-domain corpus, while [57] did not.

Reflection

Summarising related work above, our study addresses two limitations of state-of-the-art. First, despite the abundance of semantic markup data on the Web, there are only a very small number of studies that explored the use of such data to create language resources for downstream language processing tasks. Among them, the typical approach is training embedding models using such data [11, 24, 52]. However, these methods and/or resources are often ad-hoc, and their effects have not been compared on the same tasks.

Second, despite the continued interest in the research of product classification, linking, and fake review detection, the use of language resources to support such tasks has been highly inconsistent, ranging from no-use at all to using a diverse set of word-embedding models (e.g. [26, 33, 37]). It is unclear for example, if earlier success of building domain-specific LMs by continued pre-training of BERT models using in-domain corpora can be replicated in this domain. Adding to this complexity is the use of different datasets, and diverse use of machine learning models ranging from traditional algorithms (e.g. SVM, logistic regression), to deep neural networks, to pre-trained language models. The implication of this is that it is extremely difficult to compare the effect of using certain language resources on such tasks.

Motivated by these issues, our work in the following will explore three different ways of creating language resource using semantic markup data, and systematically evaluate them under uniform settings on the three different downstream tasks mentioned above.

Building Language Resources

In this section, we describe our method for the creation and evaluation of the language resources for product data mining. We begin with introducing the data sources we use to create the language resources (“Data sources”). We then discuss three different ways of using these data sources to create different types of language resources: training word-embedding models, continued pre-training of BERT-like LMs, and training machine translation models that are used as a proxy to generate product-related keywords. These language resources will be later used in the three downstream tasks, to be detailed in “Product classification”, “Product linking” and “Fake product review detection”.

Data Sources

In order to create language resources using semantic markup data for the product domain, we used the 2017 release of the structured data crawled by the WDC project. Specifically, we only downloaded and processed the class-specific subsets of the schema.org data related to `sg:Product`.⁹ This contains nearly 5 billion RDF n-quads, extracted from over 267 million web pages and over 812 thousand hosts. Each n-quad contains a subject, predicate, object, and a graph label which in this case, denotes the source URL of the n-quad.

Next, we parse this dataset to identify product offer instances, and build a Solr¹⁰ index of product offers with their attributes found in the n-quads. This is done by first searching for ‘definition n-quads’ that define a product offer instance with `http://www.w3.org/1999/02/22-rdf-syntax-ns#type` as the predicate, and either `sg:Product` or `sg:Offer` as the object (i.e. where an n-quad defines an instance of an `sg:Product` or `sg:Offer`), and then parsing other n-quads containing the same subject as the definition n-quad to create property-value pairs for each offer. Only data that are potentially English are retained. This is achieved by automatically checking if the source URL (i.e. graph label) contains a top-level domain that clearly indicates non-English websites (e.g. .fr, .cn). This Solr index is further processed to create two corpora: a product description corpus, and a product category corpus.

The *product description corpus* contains descriptions of product offers. These are extracted from the `sg:Product/description` property of each product offer. A light cleaning process is applied to ensure that only descriptions containing between 50 and 250 words are selected. This restriction is to reduce content that is likely to be very

noisy. For example, we noticed that sometimes product descriptions contain only a handful of generic words; while other times they are too long and can include the entire web page content. These texts are also normalised to keep only alpha-numeric characters. If a token contains digits only, it is replaced with a symbolic token to indicate a digit-only token. The resulting product description corpus contains over 1.9 billion tokens, extracted from over 34 million product offers.

The *product category corpus* contains over 700 thousands of product name—site-specific category pairs. These are selected from offer instances that have both an n-quad defining their name and site-specific labels. Product names are extracted from `sg:Product/name` or `sg:Offer/name` properties, while site-specific labels are extracted from `sg:Product/category` or `sg:Offer/category`. Both product names and site-specific labels are subject to a light cleaning process where only alpha-numeric characters are retained, and those containing more than 10 tokens (delimited by white space characters) or less than 2 tokens are removed. These restrictions are for the same reason—to reduce noise in the data. In addition, digit-only tokens are replaced with the same universal symbol. Further, a stop word list is used to filter out generic site-specific labels, such as `Home` and `Product`, and only pairs extracted from the top 100 largest hosts (as measured by the number of product offer instances found from each host) are kept. This is to focus on hosts that are potentially large e-commerce vendors and therefore, have defined relatively good quality site-specific categorisation schemata.

We will explain how we use these corpora to build language resources below.

Training Word-Embedding Models

The first approach to utilising the above corpora is training word embedding models. As discussed before, only a couple of studies [11, 24] used semantic markup data to train embedding models. However, [24] trained product embeddings that are ad-hoc, while our earlier work [11] developed word-embedding models that were not thoroughly evaluated. Here, following our previous work, we simply use the Gensim¹¹ implementation of the Word2Vec algorithm [65] to train word-embedding models using the product description corpus. We use the skip-gram algorithm for training, as it was shown to better represent infrequent words [65]. This fits our data well, as a notable fraction of words in our product classification and linking datasets (see Appendix 1) are not represented by the most frequent words found in the product description corpus.

⁹ http://webdatacommons.org/structureddata/2018-12/stats/schema_org_subsets.html.

¹⁰ <https://lucene.apache.org/solr>.

¹¹ <https://radimrehurek.com/gensim/>.

We use a sliding window of 10, minimum frequency threshold of 5 and text lower casing, then keeping the remaining parameters as default. The word embeddings have 300 dimensions. We refer to this as ‘product word embeddings’.

Continued Pre-training of BERT Language Models

The second approach explores the continued pre-training of large LMs. The principle of ‘continued pre-training’ of LMs has been introduced in the recent research. The idea is to take an existing LM such as BERT, and further training it on large, in-domain, unlabelled corpora (e.g. [66, 67]).

We explore the benefits of continued pre-training the BERT model on our product description corpus, and refer to the resulting LM as ‘BERT_{prod}’. Specifically, we take the ‘bert-base-uncased’ model¹² and run the masked language modelling task on our product description corpus, keeping all hyperparameters as the default.¹³

However, pre-training LMs is an extremely resource-demanding process, and due to our limited access to HPC resources, we had to split our product description corpus to small segments, and create different versions of BERT_{prod} models. Specifically, we randomly sampled 8% (or approx. 570 MB, which is the maximum size of a corpus we can fit with the pre-training process on our hardware) of our corpus for 7 times ensuring no overlap of selected product descriptions, thus creating 7 smaller corpora to continue to pre-train the BERT model. This creates 7 BERT_{prod} models, and the total size of data used for continued-pre-training represents 50% of the original product description corpus.

Training Machine Translation Models

The third approach to utilising the product corpora is inspired by the work of [13]. The authors cast product classification as an MT task, whose goal is to learn the mapping between the sequence of words in a product name, to the sequence of category labels that form a hierarchical path. In this sense, the product names and their category label paths are treated as two different languages.

However, an important difference of Li’s work from ours is that the dataset they used for training the MT models is arguably, much better quality. This is because it is collected from a single vendor website, hence there is only one categorisation scheme and the naming and categorisation of products are generally consistent. In contrary, our product category corpus contains data from hundreds of different hosts, potentially selling very different products, and therefore used

highly different and inconsistent categorisation schemata, which will have different levels of hierarchies. Further, our goal of product classification is to assign category labels from a universal schema to products from different vendors. Therefore, the site-specific categories cannot be directly used as classification targets.

Therefore, instead of using this corpus to directly train a product classifier, we use it to train MT models that map a sequence of words in the product name, to the sequence of words in the product’s site-specific category. Then given a product name in the downstream task data, we apply the MT model to generate a sequence of words, which although will unlikely to map to the end classification labels, may still be indicative of the product’s ‘type’ or ‘category’ and therefore, become useful features for the downstream tasks. We will refer to these words as ‘product related keywords’ (denoted as ‘pk’).

To train the MT model, we apply off-the-shelf MT toolkit OpenNMT [68] to the product category corpus. The encoder and decoder are 2-layer LSTM with 500 hidden units. We use all default settings of the hyperparameters in the distributed implementation.

Product Classification

In this section, we explore the usage of the different language resources created in “[Building language resources](#)” in the task of product classification. We show the datasets used for this study, and configure a number of models and compare them to evaluate the impact of these language resources on these datasets. We then present the results, which will be further discussed later in the “[Discussion](#)” section together with results from other tasks.

Datasets

We use four datasets listed in Table 1. The Rakuten dataset is the one used in the Rakuten Data Challenge [12]. This contains one million product offers crawled from Rakuten.com, an online e-commerce marketplace. Each product offer only has one type of metadata, i.e. its name. The IceCat dataset is released under the WDC project,¹⁴ and contains over 760k product offers crawled from IceCat.de, a worldwide publisher and syndicator of multilingual, standardised product data from various domains. Each offer has three types of metadata: name, description and brand. The WDC-25 dataset is also released by the WDC project,¹⁵

¹² <https://huggingface.co/bert-base-uncased>.

¹³ Using the implementation at <https://github.com/huggingface/transformers/tree/master/examples/language-modeling>.

¹⁴ <http://data.dws.informatik.uni-mannheim.de/largescaleproductcorpus/categorization/>.

¹⁵ <http://webdatacommons.org/categorization/index.html>.

Table 1 Summary of datasets for product classification

Dataset	Train	Validation	Test	Classes	Metadata
Rakuten	800k	n/a	200k	3008	Name (n)
IceCat	489,902	122,476	153,095	370	Name, description (d), brand (b)
WDC-25	20,205	n/a	4884	25	Name, description, brand, manufacturer (m)
MWPD-PC	10,012	3000	3107	lv1=37, lv2=76, lv3=281	Name, description, site-specific label/category (c)

Table 2 Configuration of input word length for neural network based classification models

Dataset	Product metadata and fixed length (tokens)
Rakuten	Name (32)
IceCat	All (256); name (32)
WDC-25	All (256); name (32)
MWPD-PC	All (256); name (32)

‘All’ refers to concatenating all product metadata detailed in Table 1 as a single text input

and contains around 24k product offers randomly sampled from over 79k websites. These are classified into a flat categorisation scheme of 25 different labels, developed with reference to the Amazon, Google and UNSPSC¹⁶ product catalogue taxonomies. This is split into a training set of over 20k offers and a test set of around 5000 offers. Each offer has a large number of metadata but only the following are selected for this work: name, description, brand, and manufacturer. The MWPD-PC dataset is the product classification dataset released in the MWPD2020 challenge [11]. It contains around 16k product offers randomly sampled from the structured product data (described by the schema.org vocabulary) crawled by the WDC project. These are classified into three levels of classification (lv1 to lv3) following the GS1 Global Product Classification standard (GPC).¹⁷ Each offer has the following metadata: name, description, and site-specific label.

The product metadata are of various word lengths from different datasets. However, neural network based classification models require text input to be of a fixed length. The normal practice is that if a real input text is shorter than this fixed length, it is padded with ‘arbitrary’ tokens. If it is longer, it is truncated. We configure the lengths according to Table 2, based on the longest input observed on the datasets and the corresponding metadata used. All the training, validation, and test splits are based on the original data releases. Our selection of datasets represents a significant

degree of diversity, containing data sourced from single vendors (Rakuten and IceCat), as well as a heterogeneous range of websites (WDC-25, MWPD-PC). Table 1 shows the statistics of these datasets.

Model Configurations

Models are configured based on the variations of the input product metadata, feature representation methods, and the machine learning algorithms. Figure 1 lists these models that will be discussed in detail below.

Using Word Embeddings

Shown in Fig. 1a (Part (a) Experiments), the baseline and their corresponding comparative models differ in terms of word-embedding representations (shaded in grey). Given a product, each model takes input all of its metadata available in a dataset, and passes them through the different word embeddings to construct a feature representation for the product. An ML algorithm then learns to classify the products based on these features.

In terms of the word embeddings that are key for comparison, we compare our product word embeddings (prod) against the generic, pre-trained Word2Vec embedding model trained using Google News (gg1).¹⁸ In terms of ML algorithms, we test a simple linear SVM (SVM), the fastText baseline used in the MWPD2020 shared task for product classification [11] (FT.MWPD), and the ‘GN-DeepCN’ structure proposed in [37] with either a biLSTM (GND.biLSTM) or HAN (GND.HAN) as its sub-structure. For SVM and fastText, different product metadata are concatenated into a single text and treated indifferently. For GND.biLSTM and GND.HAN, each specific type of product metadata is fed into a sub-structure (biLSTM or HAN) that learns separate feature representations for them. The implementation and specifications of these algorithms are as follows:

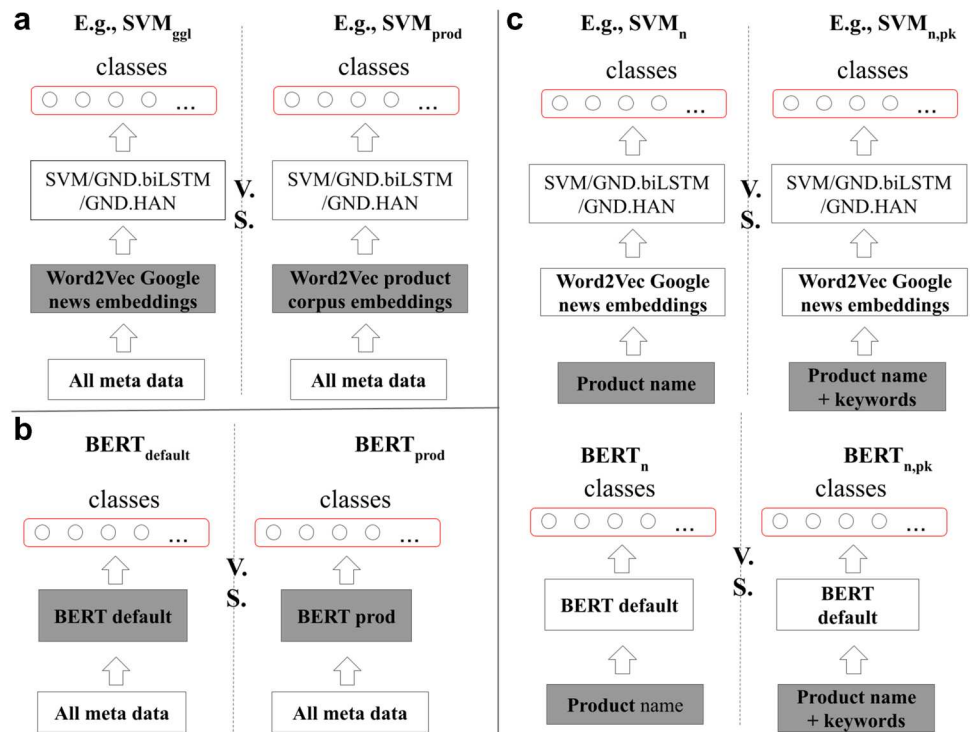
- SVM: implemented in Scikit-Learn 0.19, with the parameters set as follows: regularisation (*c*) parameter of 0.01,

¹⁶ <https://www.unspsc.org/>.

¹⁷ <https://www.gs1.org/standards/gpc>.

¹⁸ <https://code.google.com/archive/p/word2vec/>.

Fig. 1 Configurations of different models for comparison for the product classification task. The shaded box represents the components of a model to be changed for comparison



one-vs-rest multi-classification training, balanced class weight, L2 penalisation and squared hing loss.

- fastText: default implementation as in [11]
- GND.biLSTM and GND.HAN: default implementation by [37], using an epoch of 20 and a batch size of 128. All other hyperparameters remain unchanged.

Following this, a model using the product word embeddings will be compared against itself when using the generic word embeddings. E.g. SVM_{prod} against SVM_{ggl}, or GND.HAN_{prod} against GND.HAN_{ggl}.

Using Language Models

Shown in Fig. 1b (Part (b) Experiments), the baseline and their corresponding comparative models differ in terms of the underlying language models (LM) used. Given a product, each LM takes input all of its metadata available in a dataset, and passes them into the different LM (shaded in grey), which produces its feature representation and learns classification patterns in an end-to-end fashion.

In terms of the LM, we compare a generic BERT model (BERT_{default}) against the ones created following our method in “Training machine translation models” (BERT_{prod}). As mentioned before, we had to create 7 different LMs. Here, BERT_{prod} refers to the average performance recorded for all these LMs. For both BERT_{default} and BERT_{prod}, classification is achieved through stacking a linear layer on top of the corresponding language model. In the case of BERT-based

LMs, the output of the first token from the final hidden state of the model is used in the final classification. Same as SVM and FT.MWPD, product metadata are concatenated into a single piece of text. The implementation and specifications of are as follows:

- Implemented based on PyTorch 1.7.0,¹⁹ with a batch size of 32, learning rate of 2e−5, epochs of 10 and using the Adam algorithm with weight decay for optimisation. All other hyperparameters remain unchanged from the distribution.
- For BERT_{default}, the ‘bert-base-uncased’ model from the generic distribution is used.

Using Machine Translation Models

Shown in Fig. 1c (Part (c) Experiments), the baseline and their corresponding comparative models differ in terms of the product metadata used (shaded in grey). As discussed before, we apply the MT model trained in “Training machine translation models” to product names from each dataset to generate product-related keywords (pk), and use these keywords as another type of metadata for each product.

For each model described above in “Using word embeddings” and “Using language models”, we first replace the product metadata with only product names, then create two variants that are compared against each other: one that uses

¹⁹ <https://pytorch.org/>.

Table 3 Product classification results comparing the use of word-embedding models (MWPDC-PC dataset)

Product classification experiments: using word-embedding models, MWPDC-PC

Dataset	Baselines						Comparators							
	Macro			W.Macro			Macro			W.Macro				
	P	R	F1	P	R	F1	P	R	F1	P	R	F1		
MWPDC-PC (lv1)	SVM _{gg1}	61.2	61.5	60.5	80.4	80.3	80.2	SVM _{prod}	68.6	66.6	66.5	83.0	82.8	82.7
	GND.biLSTM _{gg1}	60.6	54.5	55.9	81.3	81.4	80.8	GND.biLSTM _{prod}	68.5	63.2	64.9	85.5	85.6	85.3
	GND.HAN _{gg1}	66.4	63.0	63.0	84.6	83.7	83.5	GND.HAN _{prod}	71.2	70.2	69.9	86.5	85.9	86.0
	FT.MWPD _{gg1}	73.3	68.3	69.0	86.9	86.5	86.4	FT.MWPD _{prod}	76.9	70.4	72.3	87.2	87.5	87.1
MWPDC-PC (lv2)	SVM _{gg1}	58.2	57.8	56.0	80.7	80.3	80.2	SVM _{prod}	62.4	60.2	58.2	82.6	82.2	82.0
	GND.biLSTM _{gg1}	53.7	49.9	49.4	80.6	80.6	80.1	GND.biLSTM _{prod}	60.1	55.2	56.5	85.3	85.4	85.0
	GND.HAN _{gg1}	57.2	55.0	53.5	82.5	82.2	81.6	GND.HAN _{prod}	61.6	56.7	57.0	86.4	85.9	85.5
	FT.MWPD _{gg1}	59.5	56.0	55.4	85.7	85.0	84.8	FT.MWPD _{prod}	69.4	60.9	62.3	86.9	86.8	86.4
MWPDC-PC (lv3)	SVM _{gg1}	47.4	47.1	45.2	74.1	73.6	72.8	SVM _{prod}	47.0	47.6	47.0	77.0	75.8	75.3
	GND.biLSTM _{gg1}	37.9	37.7	35.5	72.0	72.7	71.3	GND.biLSTM _{prod}	44.1	43.5	41.6	77.4	77.1	76.2
	GND.HAN _{gg1}	41.9	41.8	39.0	76.7	75.6	74.9	GND.HAN _{prod}	47.5	48.9	45.7	79.1	78.1	77.4
	FT.MWPD _{gg1}	41.2	40.1	38.5	78.6	75.8	76.1	FT.MWPD _{prod}	49.4	46.8	45.7	80.1	79.0	78.5

Boldfaced text suggests the results are better than the baseline when the language resources are used

Table 4 Product classification results comparing the use of word-embedding models (other datasets)

Product classification experiments: Using word-embedding models, other datasets

Dataset	Baselines							Comparators						
	Macro			W.Macro			Macro			W.Macro				
	P	R	F1	P	R	F1	P	R	F1	P	R	F1		
IceCat	SVM _{ggl}	92.5	95.1	93.7	99.0	98.8	98.9	SVM _{prod}	94.8	95.0	94.6	99.1	99.1	99.1
	GND.biLSTM _{ggl}	94.4	94.3	94.1	99.2	99.2	99.2	GND.biLSTM _{prod}	95.2	94.4	94.5	99.3	99.3	99.3
	GND.HAN _{ggl}	95.0	92.8	93.5	99.1	99.0	99.0	GND.HAN _{prod}	95.6	94.1	94.4	99.2	99.2	99.2
	FT.MWPD _{ggl}	93.3	94.7	93.7	99.3	99.27	99.29	FT.MWPD _{prod}	93.9	95.1	94.2	99.4	99.31	99.33
WDC-25	SVM _{ggl}	72.8	73.5	72.2	81.8	80.3	80.2	SVM _{prod}	74.3	74.9	73.9	82.7	82.0	81.8
	GND.biLSTM _{ggl}	70.6	70.2	69.5	79.1	78.5	78.1	GND.biLSTM _{prod}	75.3	76.0	74.5	84.0	82.6	82.5
	GND.HAN _{ggl}	69.1	68.3	67.7	77.9	76.7	76.5	GND.HAN _{prod}	74.3	74.1	73.1	83.2	81.2	81.4
	FT.MWPD _{ggl}	77.3	77.6	76.5	84.7	84.2	83.7	FT.MWPD _{prod}	78.2	78.6	77.5	86.0	85.4	85.0
Rakuten	SVM _{ggl}	22.5	42.8	26.5	72.1	59.4	61.2	SVM _{prod}	27.8	44.0	31.6	74.2	64.7	66.7
	GND.biLSTM _{ggl}	39.8	35.3	35.7	74.9	75.1	74.4	GND.biLSTM _{prod}	41.1	37.9	37.8	76.9	76.7	76.3
	GND.HAN _{ggl}	42.3	37.7	37.8	76.2	76.2	75.5	GND.HAN _{prod}	42.0	37.7	37.7	76.5	76.4	75.8
	FT.MWPD _{ggl}	38.6	44.7	39.9	81.8	80.6	80.9	FT.MWPD _{prod}	40.1	45.7	41.2	81.8	80.9	81.1

Boldfaced text suggests the results are better than the baseline when the language resources are used

Table 5 Product classification results comparing the use of continued pre-training of the BERT language model

Product classification experiments: continued pre-training of language models

Dataset	Baselines							Comparators						
	Macro			W.Macro			Macro			W.Macro				
	P	R	F1	P	R	F1	P	R	F1	P	R	F1		
MWPD-PC (lv1)	BERT _{default}	71.3	68.6	68.7	88.4	88.4	88.2	BERT _{prod}	73.6	71.8	71.8	89.4	89.4	89.2
MWPD-PC (lv2)	BERT _{default}	56.7	55.5	54.4	86.5	87.1	86.4	BERT _{prod}	57.3	58.1	56.6	87.3	88.2	87.4
MWPD-PC (lv3)	BERT _{default}	28.1	29.4	26.9	71.1	76.4	72.4	BERT _{prod}	31.0	32.0	29.7	73.1	78.0	74.3
IceCat	BERT _{default}	97.0	96.8	96.8	99.6	99.5	99.5	BERT _{prod}	97.2	96.8	96.8	99.6	99.6	99.6
WDC-25	BERT _{default}	80.5	79.7	79.1	86.8	85.6	85.5	BERT _{prod}	77.1	77.8	76.3	84.4	83.6	83.2
Rakuten	BERT _{default}	36.8	34.8	34.4	82.0	82.9	82.2	BERT _{prod}	36.5	34.6	34.1	81.8	82.8	82.0

Boldfaced text suggests the results are better than the baseline when the language resources are used

the product name only, the other using the name (n) plus the product-related keywords (n,pk, in this case the fixed length for input text is set to double that of name, i.e. 64). In addition, each model only uses the generic language resources (i.e. the generic word embeddings, or the BERT LM). This is to exclude the effects from all other factors, thus allowing the results to focus on the use of product related keywords.

As examples, SVM_n is compared against SVM_{n,pk}, both using the generic Google News word embeddings; while

BERT_n is compared against BERT_{n,pk}, both using the bert-base-uncased generic LM.

Evaluation Metrics

In terms of evaluation metrics, we use the standard Precision (P), Recall (R) and F1 scores for classification tasks. These are calculated using Eqs. (1)–(3), where TP denotes True Positives, FP denotes False Positives, and FN denotes False

Negatives. As some of our datasets contain highly unbalanced classes (e.g. MWPD-PC), we report macro-averages across all classes (arithmetic mean of individual classes' P, R, F1 scores) in order to analyse a classifier's performance on small classes, as well as weighted macro-averages (similar to macro-averages but weighs the score of each class label by the number of true instances when calculating the average) which was used in [11] for ranking all participating systems.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

$$\text{F1} = \left(\frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \right) \quad (3)$$

Result Summary

In terms of the effects of using word embedding models, Tables 3 and 4 show that our skip-gram word-embedding model trained on the product description corpus is able to bring consistent improvement on all datasets, with all classifiers. This improvement is noticed for Precision, Recall, and F1 (macro- and weighted macro-average), with only a handful of exceptions where the results were very close to the baseline. For example, on the Rakuten dataset, the GND.HAN_{skip-all} obtains a macro-F1 of 37.7, which is lower but still comparable to the corresponding baseline GND.HAN_{ggl}'s 37.8. The improvement can be significant in many cases, such as 9.0 in macro-F1 by GND.biLSTM_{prod} against GND.biLSTM_{ggl} on the MWPD-PC (lv11) dataset (row 5 Table 4), and 6.9 in macro-F1 by FT.MWPD_{prod} against FT.MWPD_{ggl} on the MWPD-PC (lv12) dataset (row 11 Table 3). The improvement on the IceCat is the smallest, but consistent. The baselines on this dataset already achieved very high F1.

In terms of the effects of continued pre-training of the LM, Table 5 shows less promising results. We are unable to obtain consistent improvement on all datasets, but only on MWPD-PC and the IceCat datasets, where the improvement is very small. One may argue that a potential reason for the better results on the MWPD-PC dataset is the possible similarity between the corpus used to create this gold standard, and the corpus used to continue pre-training the BERT LM. Both are based on the n-quad corpora released by the WDC project. However, we expect such impact to be minimal. On the one hand, we ensured that different releases were used (Nov 2017 release for the product description corpus, and a mixture of Nov 2018

and pre-2014 releases for the MWPD-PC goldstandard²⁰). On the other hand, the releases were based on random crawls of the Web. Interestingly, BERT-based classifiers have achieved better results than SVM, GND based structures, or FT.MWPD on all datasets except MWPD-PC lv13, which is harder due to more fine-grained classes.

In terms of the effects of MT-based product keywords, Tables 6 and 7 show that they do not bring consistent benefits, regardless of datasets or classifiers. Although there are cases where such keywords bring improvements in the results, in the majority of cases, they caused classifier accuracy to decrease. The SVM classifier is the only one that benefited in most cases from such keywords on all datasets. Nevertheless, we cannot conclude such keywords as useful for product classification task.

Product Linking

In this section, we explore the usage of the different language resources created in “[Building language resources](#)” in the task of product linking. Following a similar structure to “[Product classification](#)”, we present the datasets used, configuration of models, and their evaluation results.

Datasets

As shown in Table 8, we use a total of 9 datasets from two main sources: the WDC project and the DeepMatcher project. All datasets includes pairs of product offers and a binary label indicating if the offers match or not. The WDC project released several product linking datasets by parsing and annotating samples of the CommonCrawl corpus. These are used in later studies such as [22]. We use the ‘small’ dataset as reported in [22] for a number of reasons. First, the ‘small’, ‘medium’, ‘large’ and ‘extra large’ datasets all have the same test set. The only difference is the size of the training set, which contains a different number of instances created in a distantly supervised manner. Second, our choice is also limited by our computation resources. Each offer has the following metadata: name, description, price, brand, specification table as a text, specification key-value pairs, and site-specific label.

The DeepMatcher project released 13 datasets for evaluating entity linking while not all of them are related to the product domain. These were further split into three groups: ‘structured’ where product metadata are defined as key-value pairs, with values being atomic, i.e. short and pure, and not a composition of multiple values that should appear separately; ‘textual’ where product metadata are long textual blobs (e.g. long title, short description); and ‘dirty’ where product

²⁰ For details, see [11].

metadata are structured, but the values for some attributes could be misplaced, or empty. We selected 8 datasets that are arguably product-related, containing 5 structured datasets, 1 textual and 2 dirty datasets. Each dataset is split into train, test and validation sets with a ratio of 3:1:1.

Similar to the product classification datasets, the product metadata are of various word lengths and for neural network-based classification models, we need to set a fixed length for them when they are used as input texts. These are configured according to Table 9. All the training, validation, and test splits are based on the original data releases. Similar to the classification task, our selection of datasets is very diverse, as shown in Table 3.

Model Configurations

Again, since our focus is evaluating the effect of different language resources on this task, we use ‘out of the box’ state-of-the-art solutions to configure different models using different language resources for comparison. Specifically, we use DeepMatcher [69] and the Natural Language Inference (NLI) model based on BERT [70].

DeepMatcher (DM) is a software package²¹ implementing state-of-the-art entity linking algorithms using DNNs. It splits the matching process into three modules: the attribute embedding module transforms input textual data of an entity mention into word embedding-based representations; the similarity representation module learns a representation that captures the similarity of two entity mentions using their embedding representations; and the classifier module that takes as input the similarity representations to determine if the two entity mentions should match or not. The similarity representation module has two key components: attribute summarisation that implements different DNN structures for interpreting the embedding representations of input entities; and attribute comparison that implements different measures for comparing the ‘summary vectors’ generated by the summarisation component. In this work, we configure DeepMatcher as follows:

- Similarity representation module: we use a ‘hybrid’ attribute summariser and the ‘element-wise absolute difference’ attribute comparator, as these were found to be the optimal settings for a wide range of scenarios
- Classifier module: we use the multi-layer NN, which is the only option available
- Attribute embedding module: this is a factor for comparison and is detailed in the section below.

Other hyperparameters of DM remain unchanged from the default software distribution.

For the BERT-based NLI model (BERT), we simply use a state-of-the-art implementation by Keras.²² The model consists of two channels, each taking one sentence as input to learn a representation vector. These vectors are then concatenated and passed to a simple linear structure for classification, which determines if the two sentences entail each other or not. We simply consider each product entity as a ‘sentence’ and construct a textual representation that fits the model. All specification and configurations remain unchanged from the implementation above.

Next, using different language resources and/or product metadata, Fig. 2 lists variants of DM and BERT to be discussed in detail below.

Using Word Embeddings

Shown in Fig. 2a (Part (a) Experiments), DM using a built-in generic word-embedding model (DM_{default} baseline) is compared with DM using the product word-embedding model (DM_{prod}). Given a product, all of its metadata are concatenated into a single text input.

Using Language Models

Shown in Fig. 2b (Part (b) Experiments), the BERT NLI model either uses the default, generic LM ‘bert-base-uncased’ (BERT_{default} baseline), or the product LMs (BERT_{prod}). Same as product classification, BERT_{prod} refers to the average performance recorded for all the seven product LMs. Product metadata are also concatenated as a single text input.

Using Machine Translation Models

Shown in Fig. 5.2.2c (Part (c) Experiments), the baseline and their corresponding comparative models differ in terms of the product metadata used. Following the same way as product classification experiments, the MT model is applied to product names from each dataset to generate product related keywords (pk), which are used as another type of metadata for each product. Then DM and BERT (each using their generic word embeddings and LM respectively) will use either only product names as input (n), or product name plus product keywords as input (n,pk, in this case the fixed length for input text is set to 64).

Evaluation Metrics

Since all datasets treat the task as binary classification, we use the same evaluation metrics as product classification.

²¹ <https://github.com/anhaidgroup/deepmatcher>.

²² https://keras.io/examples/nlp/semantic_similarity_with_bert/.

Table 6 Product classification results comparing the use of MT-generated product keywords (MWPDP-PC dataset)

Product classification experiments: using machine translation models, MWPDP-PC

Dataset	Baselines						Comparators							
	Macro			W.Macro			Macro			W.Macro				
	P	R	F1	P	R	F1	P	R	F1	P	R	F1		
MWPDP-PC (lvl1)	SVM _n	54.7	63.7	57.9	79.3	77.3	77.9	SVM _{n,pk}	60.7	66.7	62.8	80.8	79.9	80.0
	GND.biLSTM _n	59.9	59.8	59.1	81.6	81.0	80.8	GND.biLSTM _{n,pk}	64.2	59.8	60.8	81.6	81.3	81.0
	GND.HAN _n	66.2	62.5	63.1	81.8	81.2	81.1	GND.HAN _{n,pk}	65.2	63.8	63.8	82.7	82.7	82.4
	FT.MWPD _n	73.6	68.1	69.5	87.0	86.8	86.6	FT.MWPD _{n,pk}	67.5	63.3	64.6	84.7	83.9	84.0
	BERT _n	68.6	67.4	67.0	86.9	87.2	86.9	BERT _{n,pk}	68.3	64.7	64.6	86.2	86.2	85.9
MWPDP-PC (lvl2)	SVM _n	57.2	67.6	59.6	79.8	77.3	78.0	SVM _{n,pk}	58.2	66.2	59.8	80.5	79.2	79.4
	GND.biLSTM _n	59.9	55.0	54.8	81.1	80.2	80.1	GND.biLSTM _{n,pk}	54.1	49.1	49.7	81.0	80.5	80.2
	GND.HAN _n	61.7	59.4	57.6	81.7	80.8	80.6	GND.HAN _{n,pk}	54.8	51.8	51.1	80.7	80.4	80.0
	FT.MWPD _n	67.4	61.7	62.4	86.0	85.5	85.3	FT.MWPD _{n,pk}	58.7	54.6	53.9	83.7	82.5	82.6
	BERT _n	49.4	48.7	47.3	84.0	84.8	84.0	BERT _{n,pk}	48.9	47.1	45.4	83.3	84.3	83.3
MWPDP-PC (lvl3)	SVM _n	50.1	59.0	51.9	77.5	73.2	74.0	SVM _{n,pk}	51.4	57.7	52.2	77.3	74.3	74.6
	GND.biLSTM _n	45.2	43.9	42.5	73.3	73.1	72.3	GND.biLSTM _{n,pk}	42.5	40.8	39.3	72.7	72.6	71.5
	GND.HAN _n	49.4	47.9	46.1	77.2	76.0	75.3	GND.HAN _{n,pk}	48.0	46.0	44.9	74.8	75.0	74.0
	FT.MWPD _n	54.3	53.5	51.6	82.7	80.6	80.8	FT.MWPD _{n,pk}	44.5	43.8	41.9	78.7	74.9	75.8
	BERT _n	23.0	26.8	23.4	66.3	73.6	68.6	BERT _{n,pk}	23.4	20.2	23.1	63.9	71.4	66.0

Boldfaced text suggests the results are better than thebaseline when the language resources are used

Table 7 Product classification results comparing the use of MT-generated product keywords (other datasets)

Product classification eperiments: using machine translation models, other datasets

Dataset	Baselines						Comparators							
	Macro			W.Macro			Macro			W.Macro				
	P	R	F1	P	R	F1	P	R	F1	P	R	F1		
IceCat	SVM _n	79.1	91.0	83.4	96.9	96.1	96.4	SVM _{n,pk}	81.6	91.0	85.1	97.1	96.5	96.7
	GND.biLSTM _n	91.3	88.2	89.1	98.1	98.0	98.0	GND.biLSTM _{n,pk}	90.7	88.4	89.0	98.2	98.1	98.1
	GND.HAN _n	90.7	86.8	88.0	98.0	97.9	97.9	GND.HAN _{n,pk}	90.9	87.4	88.4	98.1	97.9	97.92
	FT.MWPD _n	91.7	93.0	92.1	98.8	98.7	98.8	FT.MWPD _{n,pk}	89.8	91.9	90.5	98.7	98.6	98.6
	BERT _n	94.1	93.2	93.4	98.9	98.9	98.9	BERT _{n,pk}	94.1	93.2	93.4	98.9	98.9	98.9
WDC-25	SVM _n	63.7	65.5	63.3	76.3	73.6	74.1	SVM _{n,pk}	62.5	63.8	61.9	73.7	72.4	72.3
	GND.biLSTM _n	63.5	63.3	62.6	74.5	73.3	73.2	GND.biLSTM _{n,pk}	63.8	60.6	60.7	72.2	71.9	70.8
	GND.HAN _n	62.1	63.0	61.2	73.1	71.6	71.3	GND.HAN _{n,pk}	60.1	60.5	58.3	71.0	69.4	68.9
	FT.MWPD _n	69.7	69.1	67.5	79.9	78.6	77.8	FT.MWPD _{n,pk}	65.8	65.0	63.5	77.9	74.5	75.0
	BERT _n	70.6	71.0	69.5	80.8	78.6	78.7	BERT _{n,pk}	71.4	72.3	70.7	81.8	79.9	80.0
Rakuten	SVM _n	22.5	42.5	26.5	72.1	59.4	61.2	SVM _{n,pk}	24.1	41.9	28.0	74.2	67.4	66.7
	GND.biLSTM _n	39.8	35.3	35.7	74.9	75.1	74.4	GND.biLSTM _{n,pk}	38.7	34.8	35.0	74.7	75.0	74.4
	GND.HAN _n	42.3	37.7	37.8	76.2	76.2	75.5	GND.HAN _{n,pk}	39.3	34.5	34.9	74.6	74.6	73.9
	FT.MWPD _n	38.6	44.7	39.9	81.8	80.6	80.9	FT.MWPD _{n,pk}	36.1	40.8	36.8	80.8	79.3	79.8
	BERT _n	36.8	34.8	34.4	82.0	82.9	82.2	BERT _{n,pk}	36.3	34.2	33.8	81.6	82.6	81.8

Boldfaced text suggests the results are better than thebaseline when the language resources are used

Table 8 Summary of datasets for product linking

Dataset	Train	Validation	Test	Domain	Metadata
WDC-small	7230	1808	4400	Camera, computer, shoes, watches	Name, description, brand etc. 7 types
BeerAdvo-RateBeer (S)	268	91	91	Beer	Name, brewer, style, ABV
iTunes-Amazon ₁ (S)	321	109	109	Music	Name, artist, album etc. 8 types
Fodors-Zagats (S)	567	190	190	Restaurant	Name, address, city etc. 6 types
Amazon-Google (S)	6874	2293	2293	Software	Name, manufacturer, price
Walmart-Amazon ₁ (S)	6144	2049	2049	Electronics	Name, category, brand etc. 5 types
Abt-Buy (T)	5743	1916	1916	Product	Name, description, price
iTunes-Amazon ₂ (D)	321	109	109	Music	Same as iTunes-Amazon ₁ but misplaced
Walmart-Amazon ₂ (D)	6874	2293	2293	Software	Same as Walmart-Amazon ₁ but misplaced

S structured, T textual, D dirty

The only difference is that, same as the literature, these are computed for true positives only (i.e. true matches).

Result Summary

In terms of the word-embedding model (Table 10), our skip-gram-based word embeddings can further improve F1 on six out of nine datasets. However, on the other three datasets they caused significant decline in F1. Referring to Table 3, we would argue that the three datasets where the decline is noted may be either too small (BeerAdvo-RateBeer (S)), or less relevant to the conventional ‘product’ domain (Fodors-Zagats (S), restaurant; Amazon-Google (S), software).

In terms of the continued pre-training of LM and MT-based product keywords, based on results in Tables 11 and 12, we are unable to conclude either to be useful for this task. Positive improvements can be noticed on some datasets, but these have been very inconsistent.

Fake Product Review Detection

While the previous two tasks concern data that are typically properties of products, fake product review detection concerns data that is rather indirectly connected to products. For this task, we only experiment with the use of word-embedding models and in-domain LMs, not the MT model. This is because the typical review datasets do not contain product names which we require as input to the MT model. Using the review text as input will not make sense because the MT model is trained to learn mappings between short sequences of words. In addition, the vocabularies used during training are very different.

Datasets

We use the dataset from [57], which were created using a Natural Language Generation model and contains over 40,000 reviews of products from 10 broad categories (e.g. Books, Electronics), each labelled as either a fake review, or genuine. For neural network-based classifiers that require a fixed input text length, this is set to 512. Although the reviews are automatically generated by algorithms, the authors showed that these proved to be harder for human annotators to differentiate. We do not expand our experiments to other fake review datasets, because as we shall show later, we have observed same patterns as those in the other two tasks and we do not expect adding more datasets to add further values to our findings.

Table 9 Configuration of input word length for neural network-based models. ‘all’ refers to concatenating all product metadata detailed in Table 3 as a single text input

Dataset	Product metadata and fixed length (tokens)
WDC-small	All (512); n (32)
BeerAdvo-RateBeer (S)	All (64); n (32)
iTunes-Amazon ₁	All (128); n (32)
Fodors-Zagats (S)	All (64); n (32)
Amazon-Google (S)	All (64); n (32)
Walmart-Amazon ₁ (S)	All (64); n (32)
Abt-Buy (T)	All (128); n (32)
iTunes-Amazon ₂ (D)	All (128); n (32)
Walmart-Amazon ₂ (D)	All (64); n (32)

Model Configurations and Evaluation Metrics

Since fake review detection is a binary text classification task, our model configurations will follow those from the product classification task (“[Model configurations](#)”). The dataset only has one source of text input, namely, the review text. Therefore, the models only vary in terms of the underlying language resources used. Figure 3 lists these models that will be briefly covered below.

In terms of using word-embedding models, shown in Fig. 3a (Part (a) Experiments), the baseline and their corresponding comparative models differ in terms of word embedding

representations (shaded in grey). Same as product classification, we compare our product word embeddings (prod) against the generic, pre-trained Word2Vec embedding model (ggl). We combine all algorithms listed in “[Using word embeddings](#)” with each of the two options of word-embedding models. The same configuration and specifications are used.

In terms of using in-domain LMs, shown in Fig. 3b (Part (b) Experiments), the baseline and their corresponding comparative models differ in terms of the underlying LMs used. Again we use the same model variants from “[Using language models](#)”, but on the different dataset. The same configuration and specifications are retained.

In terms of evaluation, the same metrics for product classification explained in “[Evaluation Metrics](#)” are used here.

Result Summary

Overall(see Table 13), we observe same patterns as product classification and linking tasks. On the one hand, product word embeddings trained on the product description corpus led to consistent improvement in F1 on this task over the generic word-embedding model, with the highest noted with the SVM classifier and the lowest noted with the fastText classifier. Compared to the product classification task, it is worth highlighting that the text content in the two datasets can be notably different. This suggests that the ‘knowledge’

Fig. 2 Configurations of different models for comparison for the product linking task. The shaded box represents the components of a model to be changed for comparison

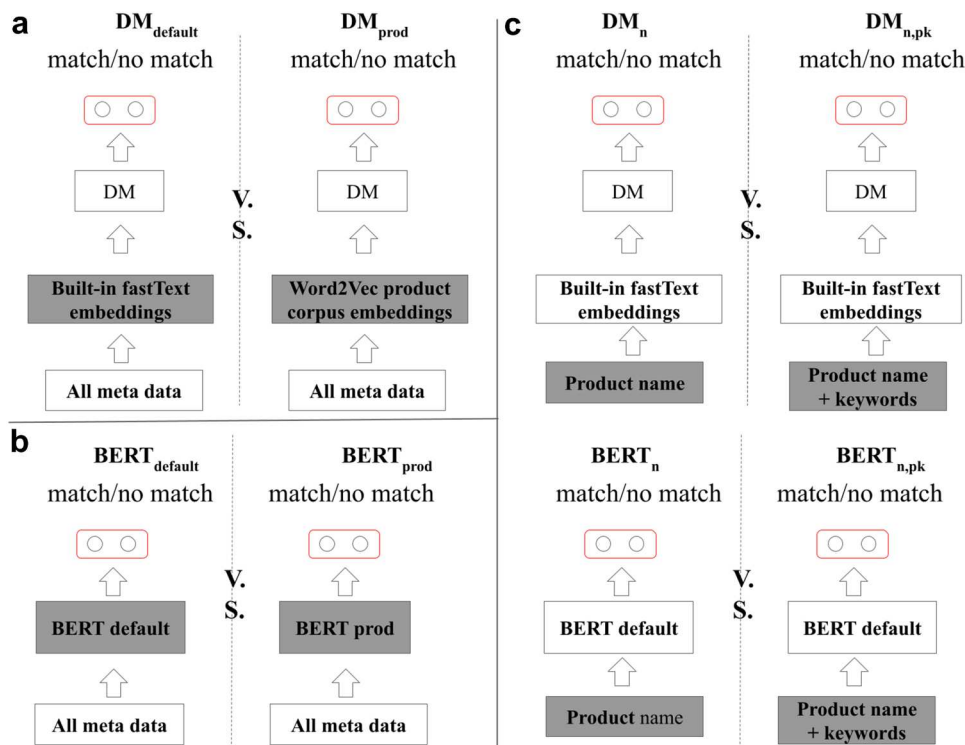


Table 10 Product linking results comparing the use of word embedding models

Product linking experiment: using word-embedding models						
Dataset	Baseline (DM _{default})			Comparator (DM _{prod})		
	P	R	F1	P	R	F1
WDC-small	57.3	63.5	60.3	66.2	65.2	65.7
BeerAdvo-RateBeer (S)	66.7	71.4	69.0	57.1	57.1	57.1
iTunes-Amazon ₁ (S)	86.2	92.6	89.3	96.0	88.9	92.3
Fodors-Zagats (S)	100.0	95.5	96.7	90.9	90.9	90.9
Amazon-Google (S)	58.8	60.3	59.5	51.0	41.9	46.0
Walmart-Amazon ₁ (S)	33.9	20.2	25.3	40.5	23.3	29.6
Abt-Buy (T)	26.7	18.9	22.2	25.5	23.3	24.4
iTunes-Amazon ₂ (D)	57.6	70.4	63.3	72.7	59.3	65.3
Walmart-Amazon ₂ (D)	23.6	17.1	19.8	28.3	24.4	26.2

Boldfaced text suggests the results are better than the baseline when the language resources are used

Table 11 Product linking results comparing the use of language models

Product linking experiments: using language models						
Dataset	Baseline			Comparator		
	(BERT _{default})			(BERT _{prod})		
	P	R	F1	P	R	F1
WDC-small	64.2	83.5	72.3	63.2	85.0	72.1
BeerAdvo-RateBeer (S)	63.3	91.4	73.7	62.5	88.4	71.1
iTunes-Amazon ₁ (S)	82.5	86.7	84.4	81.6	85.0	82.9
Fodors-Zagats (S)	89.0	99.1	93.6	88.7	97.8	92.7
Amazon-Google (S)	68.6	68.6	68.0	64.3	70.5	66.3
Walmart-Amazon ₁ (S)	66.3	76.0	70.5	58.9	72.4	64.4
Abt-Buy (T)	83.4	69.0	75.3	74.3	72.8	72.9
iTunes-Amazon ₂ (D)	87.0	81.5	84.0	81.5	80.4	80.4
Walmart-Amazon ₂ (D)	66.3	69.2	67.6	55.2	71.6	61.2

Boldfaced text suggests the results are better than the baseline when the language resources are used

captured by the product word embeddings can be potentially transferable to more general product data mining tasks. On the other hand, continued pre-training of BERT still led to detrimental effects.

Further Analysis

In this section, we conduct further analysis of our datasets in order to better understand the potential contributing factors to the overall negative results.

Data Provenance

One potential cause of a less good quality training dataset is unbalanced data distribution. To understand if this could have been an issue in our study, we analysed the dominating

hosts that contributed to the product description and the product category corpora in order to discover if there exist certain dominating hosts selling a restricted range of products. To do this, we manually inspected the top 100 largest hosts as measured by the number of product offer instances found from each host, and classified them based on the types of products they sell.

As Table 14 shows, a significant portion of the dominating hosts sell fashion related products, typically clothing, footwear, accessories and so on. Therefore we expect that our product description and category corpora to contain a significant portion of data related to these domains. Notice that products from other domains such as software, beer, and restaurant are under-represented, which may help explain the observation that our word embeddings were not useful on the three datasets mentioned before in the product linking task. However, this analysis does not help explain why other language resources, i.e.e the language

Table 12 Product linking results comparing the use of machine translation models for generating product keywords

Product linking experiments: using machine translation models								
Dataset	Baselines			Comparators				
		P	R	F1		P	R	F1
WDC-small	DM _n	55.6	59.9	57.7	DM _{n,pk}	51.3	55.5	53.3
	BERT _n	67.1	89.2	76.4	BERT _{n,pk}	66.5	87.1	75.3
BeerAdvo-RateBeer (S)	DM _n	75.0	21.4	33.3	DM _{n,pk}	83.3	35.7	50.0
	BERT _n	76.2	68.6	68.8	BERT _{n,pk}	60.0	70.0	64.4
iTunes-Amazon ₁ (S)	DM _n	86.2	92.6	89.3	DM _{n,pk}	100.0	85.2	92.0
	BERT _n	86.9	92.6	89.5	BERT _{n,pk}	87.8	88.9	88.2
Fodors-Zagats (S)	DM _n	80.0	72.7	76.2	DM _{n,pk}	76.2	72.7	74.4
	BERT _n	91.3	88.2	89.5	BERT _{n,pk}	87.6	90.0	88.8
Amazon-Google (S)	DM _n	59.5	57.7	58.6	DM _{n,pk}	58.5	54.7	56.5
	BERT _n	62.5	69.3	64.9	BERT _{n,pk}	57.7	64.8	60.6
Walmart-Amazon ₁	DM _n	36.6	36.8	36.7	DM _{n,pk}	37.2	18.1	24.4
	BERT _n	61.5	55.2	56.5	BERT _{n,pk}	63.0	49.7	54.8
Abt-Buy (T)	DM _n	33.3	27.2	30.0	DM _{n,pk}	33.2	30.1	31.2
	BERT _n	76.3	79.8	77.9	BERT _{n,pk}	81.9	75.6	78.4
iTunes-Amazon ₂ (D)	DM _n	63.6	51.9	57.1	DM _{n,pk}	70.0	51.9	59.6
	BERT _n	81.9	69.6	74.3	BERT _{n,pk}	87.7	74.8	80.4
Walmart-Amazon ₂	DM _n	34.7	34.7	34.7	DM _{n,pk}	28.2	24.9	26.5
	BERT _n	62.5	55.4	56.8	BERT _{n,pk}	64.6	51.6	56.1

Boldfaced text suggests the results are better than the baseline when the language resources are used

model and MT-based product keywords are not as useful as word embeddings.

Vocabulary Coverage

Here, we study the extent to which the vocabulary of the corpus we used to build the language resources well-represents those from the target tasks. Taking the product classification and linking datasets as examples, we calculate a number of statistics on the training set of each dataset, and show them in Table 15. Avg tok is the average number of tokens (separated by the white space character) per instance (concatenating all metadata available) within a dataset. Avg % non-digit tok is the ratio between the average number of tokens excluding tokens containing only digits and the average number of tokens per instance. Avg % non-digit toks in PDC is the ratio between the average number of non-digit tokens found in the vocabulary of the product description corpus (PDC) and the average number of tokens per instance. In other words, Avg % toks in PDC indicates how much training data are covered by the vocabulary of the product description corpus.

Comparing the product classification datasets against the linking datasets, there is obvious pattern that the product linking datasets contain a much larger percentage of non-alphabetic tokens (Avg % non-digit tok), and a much smaller percentage of alphabetic tokens are covered by the product

description corpus (Avg % non-digit toks in PDC). We argue that this difference could explain why our word embeddings and fine-tuned BERT language model are found to be less effective on the product linking task.

Among the product linking dataset, recall that in Table 11 showing the results of using word embeddings, BeerAdvo-RateBeer (S), Fodors-Zagats (S) and Amazon-Google (S) are the three datasets where our word-embedding model caused the baseline performance to decline. Inspecting these datasets in Table 15, we find them to share the following three patterns: the text content is short (Avg tok), having a large percentage of non-alphabetic tokens (e.g. 58% for Fodors-Zagats (S), 29% for BeerAdvo-RateBeer), and having a relatively small percentage of vocabulary coverage by the product description corpus (e.g. Avg % non-digit toks in PDC for Fodors-Zagats (S) is only 43%). The extreme example is the Fodors-Zagats (S) dataset. Each instance contains an average of 19 tokens, where only 8 are alphabetic words, and out of which only 3 or 4 are covered by the product description corpus. As a result, the learning algorithms could have been very sensitive to those very few number of words which are covered by the vocabulary. Adding to this the potential problem of under-representation of these domains as discussed in the previous section, the combination of these factors could explain the decline in performance when those word embeddings are used.

Fig. 3 Configurations of different models for comparison for the fake product review detection task. The shaded box represents the components of a model to be changed for comparison

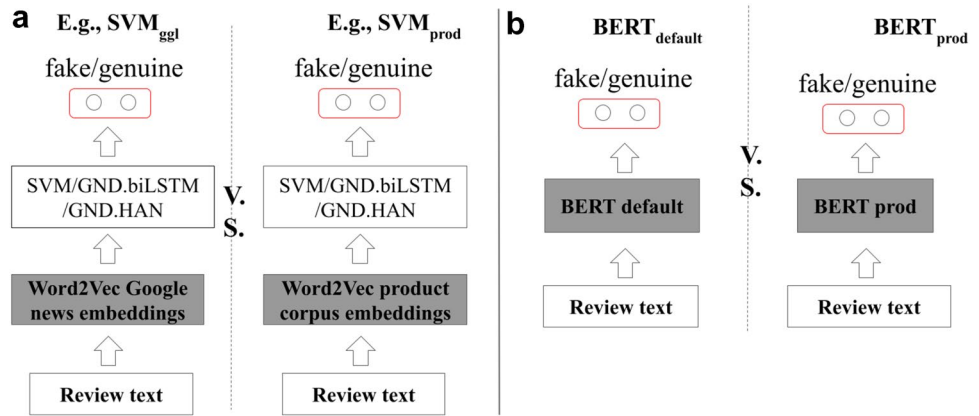


Table 13 Fake product review detection results comparing the use of word-embedding models and language models

	Macro			W.Macro		
	P	R	F1	P	R	F1
Part (a) Experiments: using word-embedding models						
Baselines						
SVM _{ggl}	85.1	85.1	85.1	85.1	85.1	85.1
GND.biLSTM _{ggl}	94.8	94.7	94.7	94.8	94.7	94.7
GND.HAN _{ggl}	95.0	95.0	95.0	95.0	95.0	95.0
FT.MWPD _{ggl}	93.9	93.9	93.9	93.9	93.9	93.9
Comparators						
SVM _{prod}	87.1	87.0	87.0	87.1	87.0	87.0
GND.biLSTM _{prod}	95.5	95.5	95.4	95.5	95.5	95.4
GND.HAN _{prod}	95.8	95.8	95.8	95.8	95.8	95.8
FT.MWPD _{prod}	94.1	94.1	94.1	94.1	94.1	94.1
Part (b) Experiments: using language models						
Baseline						
BERT _{default}	97.3	97.2	97.2	97.3	97.2	97.2
Comparator						
BERT _{prod}	97.1	97.0	97.0	97.1	97.0	97.0

Boldfaced text suggests the results are better than the baseline when the language resources are used

Product Keywords Analysis

Here, we focus on understanding the failure of the MT-generated product keywords. Our idea of MT-based product keywords is inspired by the work of [13], who cast product classification as a MT task that aims to learn the mapping between product names and their category classes. They successfully tested this approach on the Rakuten dataset. Therefore we compare the Rakuten dataset against the product category corpus (PCC) to understand if there exists any difference between the two datasets that have been used to train MT models. For each instance in a dataset, we count the number of tokens in its product name, and also the number of tokens in the classification label. For the Rakuten dataset, product classifications are pseudonymised by an ID number, such as 1608 > 2320 > 2173 > 2878. Li et al. [13] treated

them as a sequence of tokens. Therefore, we count the number of tokens separated by ‘>’. For the PCC, the equivalent product classifications are those site-specific product categories/labels, which we used to train the MT model. We compare the distribution of the numbers of words in Fig. 4. Further, we also count the number of unique tokens found in the product name and classification/site-specific categories from each dataset, and calculate a ratio as ‘name/category unique word ratio’. This number is 116 for Rakuten, and 4 for PCC.

Figure 4 shows that while the numbers of words in the product classifications from both datasets appear to be generally comparable, the Rakuten dataset however, contains generally longer product names than the PCC. The name/category unique word ratio of Rakuten is also orders of magnitude higher than that of the PCC. We say that the PCC as

Table 14 The numbers of the largest hosts (top 100 ranked by the number of product offer instances found in product description and category corpora) by types of products sold online

Type	#Hosts	Notes
Auto parts/accessories	5	E.g. car parts, car audio
Automobile	7	E.g. cars, motorcycles
B2B marketplace for exporter/importers	1	
Books	2	Incl. conventional and audio books
Business catalogues	1	E.g. yellow pages
Camera	1	
Chemical products	1	
Educational resources	1	E.g., posters, exercise books
Farming equipment	1	E.g. tractors, seeds, clothing
Fashion	26	E.g. clothing, footwear, accessories
Finance	1	E.g. credit card shopping
Food	1	
Gardening	1	Incl. equipment, plants, decoration etc.
Healthy supplements	1	
Hobbies/handcrafting	1	E.g. knitting supplies
Holiday making	4	E.g. hotels, flights, package holidays
Home furnishing/furniture	4	
Information consultancy	4	E.g. news, drug patent consultancy, DIY
Jewellery/watch	7	
Lighting equipment	1	
Music	1	
Office supply	2	
Pet supply	1	
Power tools	1	
E-commerce integration/comparison platforms	5	E.g. groupon, gumtree
Properties	2	
Restaurant equipment/hardware supply	2	
Speciality clothing	7	All hosts are for bridal wear
Sports equipment	4	E.g. golf, baseball
Sports fan shop	2	E.g. baseball team fan shops selling anything ranging from clothing to decoration items
Visual content	4	Typically online photo/gallery shops, e.g. vectorstock.com

Note that the number will add up above 100, as there are several websites classified under multiple product types

training data for MT is much ‘sparser’ than the Rakuten dataset. Intuitively, it would have been easier to generalise on the Rakuten dataset, as a significantly larger number of tokens in the product names would be mapped to the same token in the product classification. On the contrary, there are significantly fewer examples to learn this mapping on the PCC.

We conclude here that the reason behind the failure of the MT-based product keywords is the inconsistent quality of the generated keywords. And this could have been due to the sparsity in the product category corpus that we used to train the MT model.

Discussion

With the growing amount and availability of semantic markup data on the Web, research has started looking at how such a gigantic data resource can be used to support various data mining tasks [4, 5, 11]. However, there has been a significant variation in terms of the data sources used, the tasks addressed, the processes applied, and the findings reported. In this section, we discuss our work from four perspectives: (1) research question 1; (2) research question 2; and (3) how to interpret the generally negative results of pre-training language models.

Table 15 Vocabulary analysis of each dataset

Dataset	Avg tok	Avg % non-digit tok	Avg % non-digit toks in PDC
MWPD-PC	80	94	88
WDC-25	49	98	94
Rakuten	10	91	84
IceCat	46	99	85
WDC-small	12	83	79
BeerAdvo-RateBeer (S)	17	71	71
iTunes-Amazon ₁ (S)	41	63	59
Fodors-Zagats (S)	19	42	43
Amazon-Google (S)	11	64	63
Walmart-Amazon ₁ (S)	17	76	66
Abt-Buy (T)	35	80	74
iTunes-Amazon ₂ (D)	45	64	63
Walmart-Amazon ₂ (D)	19	74	69

In terms of 1) our first research question, (also objective 1), we processed the markup data from the 2017 WDC release and extracted all n-quads related to products. Then following the work of [6, 11, 13], we transformed the n-quads to three different types of language resources respectively: word embeddings, a refined BERT LM, and a machine translation model for generating product keywords given their names. Our choice of methods represent the most popular options in product NLP, but to the best of our knowledge, our work is the first that utilises and compares these different methods in a single study.

In terms of 2) our second research question (also objective 2), we conducted a wide range of experiments including three tasks and 10 datasets of different sizes, that are used to evaluate different SoTA models using the created language resources. In terms of the scope of our experimental studies, to the best of our knowledge, there were no previous studies of a comparable scale. On the whole, we found that only word embeddings led to the most consistent improvement over all tasks. In this direction, we noted before that the earlier work by [22] used the fastText algorithm to train word embeddings using a product-related corpus collected from the semantic markup data (so called ‘self-trained embeddings’). They showed that this domain-specific embedding model marginally increased the performance of product linking on some product categories, but overall did not offer significant value. In comparison, our Word2Vec skip-gram word-embedding model gained notable improvement over the generic embedding model on the WDC-small dataset (Table 11), which was also used by [22]. To explain this difference, we suspect that the size of the training corpora and the different pre-processing in the two studies could have been the reason. The corpus used in [22] focussed on

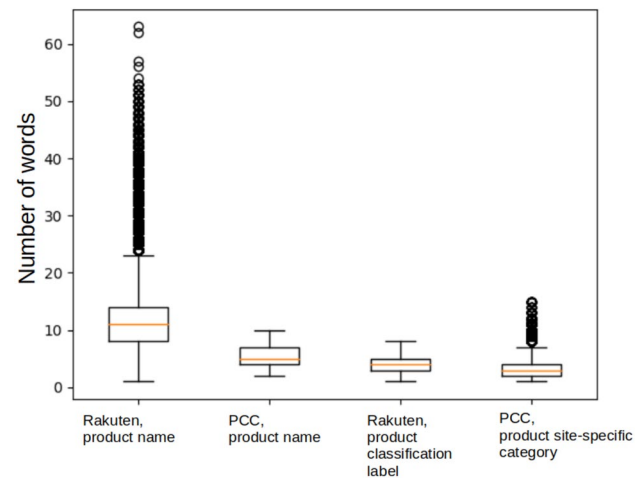


Fig. 4 Comparison of the distributions of text length in the product names from the Rakuten dataset and that from the product category corpus (PCC); and comparison of the distributions of word frequency in the product classification labels from the Rakuten dataset and the word frequency in the site-specific categories of PCC

product linking, and therefore, filtered the underlying data based on if a product offer contains useful product identifies. This process could have eliminated a significant proportion of the data that may have been useful as less than 10% of product offers contain such information. Our product description corpus on the other hand, is much larger. This seems to suggest that using a larger, more diverse set of product markup data is more beneficial for training word embeddings.

In terms of 3) the negative results from pre-training language models, we notice that, although this is generally inconsistent with the wider literature on training in-domain LMs [6–10], similar observations were also reported previously in the literature [6, 7]. We believe there can be many reasons to this, but we speculate that the primary ones being the size and quality of the data used for in-domain pre-training. Compared to the E-BERT model [9] that is the most similar to ours, we note significant difference in terms of the pre-training process. While we used the ‘out of the box’ BERT configuration without much change, E-BERT modified the pre-training process in many ways using high quality external resources and complex process algorithms. All these modifications allowed E-BERT to learn product related knowledge in a more effective way. Similarly, SMedBERT [10] changed the pre-training process by incorporating structured information such as knowledge graphs.

At the same time, it is worth noting studies that also used a simple ‘out of the box’ BERT pre-training process (same to ours) with an in-domain corpus and obtained better results on downstream tasks, such as BioBERT [6], SciBERT [8], and Clinical BERT [7]. It is possible that the main differentiating factor could be the size and

quality of the underlying corpora used for pre-training. All these earlier studies used resources that are arguably higher quality and are of a larger quantity. For example, SMedBERT, E-BERT and Clinical BERT used well-curated vocabularies, knowledge graphs or corpora. SciBERT and BioBERT used scientific publications that are well-written and follow a generally consistent structure. The unstructured in-domain corpus is typically of a comparable size to the original corpus used for training BERT, or much larger (BioBERT). In contrast, our corpora are much noisier as they are collected from heterogeneous websites and there are no standardisations on how the content should be written. Our corpora are also much smaller, due to having to split the entire dataset into smaller chunks to meet the computational restrictions.

One question that remains unanswered is that while both our word embedding model and the BERT LM are trained on the same corpus, why are the word embeddings more useful than the BERT model? On the one hand, the corpus used to pre-train BERT is much smaller, as we were unable to use the entire product description corpus as we did for the word-embedding model. On the other hand, our word embeddings are trained using the Word2Vec skip-gram algorithm, which learns word embeddings through a task of predicting the context of a given word. The BERT LM pre-training followed the Masked Language Modelling (MLM) task, which tries to predict a word given its context. This rationale is similar to Word2Vec's Continuous Bag-of-Words (CBOW) algorithm that was shown to be less effective on modelling infrequent words [65]. As Appendix 1 shows, some of the tasks we evaluated may contain many words that are under-represented in the product description corpus. As a result. This might have had an impact on the continued in-domain pre-training of BERT.

Beyond these analysis from a theoretical point of view, we also conducted further analysis of the data quality (objective 3) and this to some extent, confirmed several points made in the above discussion. We noticed that the semantic markup data are highly unbalanced in terms of the domains, with data related to fashion dominating a significant percentage. This could have led to noticeable vocabulary 'gaps', which could be detrimental to datasets containing short texts falling into such 'gaps', or algorithms that are sensitive to 'infrequent' words. Although previous work such as [4, 5] reported similar findings, they did not evaluate the impact of such problems on tasks that use such data. Our work therefore, serves first evidence of how such data quality issues impact NLP tasks.

Conclusion

In this work, we conducted an exploratory study of using structured linked data embedded within HTML webpages for the creation of language resources for downstream NLP tasks. Despite the generally negative results, we can draw important lessons that may inform future research and practice (objective 4).

From a theoretical point of view, our results serve a lesson to researchers looking to develop novel methods that exploit the growing semantic markup data on the Web. While the data has reached a 'critical mess', the sheer size does not seem to outweigh certain 'quality' aspects of the data and this may impact on downstream tasks that exploit such data. To address this, a useful task would be developing an approach to 'assess' the quality a semantic markup dataset, or identify a 'quality subset'. However, there can be many challenges such as the notion of 'quality' may be task and data dependent.

Leading from this, we argue there is also an implication on the increasingly popular research on developing large domain-specific LMs. While many successful studies have been reported in different domains, our study shows that language modelling using very large unstructured corpora may not be as straightforward as the literature indicates. Crucially, there is a lack of understanding of the 'conditions for success', e.g.: how much data would be sufficient for training a domain-specific LM, how 'balanced' the data needs to be, and in what ways? We believe that this is an important question to be further investigated, given the increasing popularity and importance of using domain-specific LMs in NLP.

From a practical point of view, our study shows that, pre-training BERT appears to be more susceptible to noise and size of datasets, while training word embeddings appears to be more robust. This can inform practitioners when making a choice between these popular approaches: although BERT-based methods are taking the mainstream in research, earlier methods like training word embeddings may perform just as well or even better in certain scenarios.

Further, we also call for further effort from data publishers that adopt the semantic markup practice. Although there has been remarkable progress in terms of the quantity of semantic markup data, it may now be the time to place more emphasis on quality in order to allow a wider community to benefit from such data.

Our work, however, is limited in a number of ways. First, partly due to the limited space of this article, we did not compare all available options of each method for creating language resources. For example, instead of Word2Vec, there are other alternatives for training word

embeddings [71]. The same can be said for pre-training LMs. Second, due to limited computing resources, we were unable to pre-train our BERT LM using the full product description corpus. This could have affected our results to some extent.

Reflecting on the above points, we highlight a few future research directions. First, as already discussed above, there can be significant value in researching the quality aspects of semantic markup data, and ultimately developing metrics and processes to identify and select the right subset of the dataset optimised for different tasks. This could be broadly considered an issue of quality of such linked data, but many research questions arise: how to define such quality metrics, how generic/task-specific can they be, how to use them to guide the dataset sub-selection, and what impact will it have on the language resources created using this dataset, and the downstream tasks using such resources?

Second, it may be worth to explore the use of these structured data in less domain-agnostic tasks. As an example, structured data embedded within specific HTML elements could be used as annotations on that web page, and the corresponding HTML formatting properties may be useful and more generalisable features for automatically tagging content from different web pages but formatted with similar properties. The idea here is how different types of content are formatted ‘relative to each other’ on a product listing web page can be consistent across many different domains and websites. This has been validated in different contexts such as [72]. The abundance of already annotated product listing pages in the form of semantic markup data can create an opportunity to train such taggers in a self-supervised way.

Our future work will explore some of the above questions and research directions.

Appendix 1: Additional Data Analysis

A.1 Product Classification: Word Frequency Analysis for Word-Embedding Model Training

As discussed in “[Training word-embedding models](#)”, we choose to use the Word2Vec skip-gram model instead of the continuous bag-of-words model for training the word-embedding models from our product description corpus. The reason is that the skip-gram model is shown to better represent infrequent words in the training corpus. We conducted a word frequency analysis of the training datasets for product classification, and discovered that a fair percentage of words belong to the relatively infrequent segment of words in the product description corpus. We present this part of analysis here.

First, we start with extracting and normalising (same process as that used for building the word-embedding model) all unique words from the product description corpus. Second, we count the frequency of these words and rank them in the descending order of frequency. This list of words is then binned into 100 segments. Third, we extract unique words from each product classification training dataset (from the concatenation of all product metadata), and count for each bin, the number of words found in that bin. Finally, we calculate the percentage of words belonging to each bin.

Our results show that, the highest frequency bin (#1) contains roughly 71%, 72%, 55% and 63% of words found in the IceCat, Rakuten, MWPD-PC and WDC-25 training sets. In other words, between 28% and 45% of words from these datasets are not the most frequent words found in the word-embedding training corpus. In fact, summing up the words under bin #11 and further, these are 3%, 6%, 21% and 17% for IceCat, Rakuten, MWPD-PC and WDC-25 training sets. Based on these findings, we opted for using the skip-gram model for training word embeddings.

Author Contributions Ziqi Zhang contributed 80% of the work (design, development, and writing); Xingyi Song contributed 20%.

Funding This work is not funded.

Availability of data and materials <https://drive.google.com/drive/folders/1BgA4iedPOFtjAeku7c4i1c2TUu7AAHp?usp=sharing>

Code availability Yes, upon request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bizer C. The emerging web of linked data. *IEEE Intell Syst.* 2009;24(5):87–92. <https://doi.org/10.1109/MIS.2009.102>.
- Foley J, Bendersky M, Josifovski V. Learning to extract local events from the web. In: Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval. SIGIR '15. Association for Computing Machinery, New York; 2015. p. 423–432. <https://doi.org/10.1145/2766462.2767739>.
- Vagliano I, Monti D, Scherp A, Morisio M. Content recommendation through semantic annotation of user reviews and linked data. In: Proceedings of the knowledge capture conference. K-CAP 2017. Association for Computing Machinery, New York; 2017. <https://doi.org/10.1145/3148011.3148035>.
- Bizer C, Primpeli A, Peeters R. Using the semantic web as a source of training data. *Datenbank Spektrum.* 2019;19:127–35.
- Meusel R, Primpeli A, Meilicke C, Paulheim H, Bizer C. Exploiting microdata annotations to consistently categorize product offers at web scale. In: International conference on electronic commerce and web technologies. Springer; 2015. p. 83–99.
- Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, Kang J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics.* 2019;36(4):1234–40. <https://doi.org/10.1093/bioinformatics/btz682>.
- Alsentzer E, Murphy J, Boag W, Weng W-H, Jindi D, Naumann T, McDermott M. Publicly available clinical BERT embeddings. In: Proceedings of the 2nd clinical natural language processing workshop. Association for Computational Linguistics, Minneapolis, Minnesota; 2019. p. 72–78. <https://doi.org/10.18653/v1/W19-1909>.
- Beltagy I, Lo K, Cohan A. SciBERT: a pretrained language model for scientific text. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP). Association for Computational Linguistics, Hong Kong, China; 2019. p. 3615–3620. <https://doi.org/10.18653/v1/D19-1371>.
- Zhang D, Yuan Z, Liu Y, Fu Z, Zhuang F, Wang P, Chen H, Xiong H. E-BERT: a phrase and product knowledge enhanced language model for e-commerce. *CoRR.* 2020. [arXiv:2009.02835](https://arxiv.org/abs/2009.02835).
- Zhang T, Cai Z, Wang C, Qiu M, Yang B, He X. SMedBERT: a knowledge-enhanced pre-trained language model with structured semantics for medical text mining. In: Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: Long Papers). Association for Computational Linguistics, Online; 2021. p. 5882–5893. <https://doi.org/10.18653/v1/2021.acl-long.457>. <https://aclanthology.org/2021.acl-long.457>
- Zhang Z, Bizer C, Peeters R, Primpeli A. Mwpa2020: semantic web challenge on mining the web of html-embedded product data. In: Proceedings of the semantic web challenge on mining the web of HTML-embedded product data co-located with the 19th international semantic web conference, vol. 2720. CEUR-WS.org; 2020.
- Lin Y-C, Das P, Datta A. Overview of the sigir 2018 ecom rakuten data challenge. In: eCOM at the 41st international ACM SIGIR conference on research and development in information retrieval. CEUR-WS.org; 2018.
- Li MY, Kok S, Tan L. Don't classify, translate: multi-level e-commerce product categorization via machine translation. *CoRR.* 2018. [arXiv:1812.05774](https://arxiv.org/abs/1812.05774).
- Yun Y, Ma D, Yang M. Human–computer interaction-based decision support system with applications in data mining. *Future Gener Comput Syst.* 2021;114:285–9. <https://doi.org/10.1016/j.future.2020.07.048>.
- Saura JR, Palacios-Marqués D, Ribeiro-Soriano D. Using data mining techniques to explore security issues in smart living environments in twitter. *Comput Commun.* 2021;179:285–95. <https://doi.org/10.1016/j.comcom.2021.08.021>.
- Saura JR, Palacios-Marqués D, Ribeiro-Soriano D. Exploring the boundaries of open innovation: evidence from social media mining. *Technovation.* 2022. <https://doi.org/10.1016/j.technovation.2021.102447>.
- Mountantonakis M, Tzitzikas Y, et al. Lodsynesis: entity extraction from text and enrichment using hundreds of linked datasets. Lecture notes in computer science. In: Harth A, et al., editors. The semantic web: ESWC 2020 satellite events. Berlin: Springer; 2020.
- Rama-Maneiro E, Vidal JC, Lama M. Collective disambiguation in entity linking based on topic coherence in semantic graphs. *Knowl Based Syst.* 2020. <https://doi.org/10.1016/j.knsys.2020.105967>.
- Kiran Selvam R, Kejriwal M. On using product-specific schema.org from web data commons: an empirical set of best practices. *arXiv e-prints.* 2020.
- Chortaras A, Stamou G. D2rml: integrating heterogeneous data and web services into custom rdf graphs. In: Workshop on Linked Data on the Web. 2018.
- Primpeli A, Peeters R, Bizer C. The wdc training dataset and gold standard for large-scale product matching. In: Companion proceedings of the 2019 world wide web conference. WWW '19. Association for Computing Machinery, New York; 2019. p. 381–386. <https://doi.org/10.1145/3308560.3316609>.
- Peeters R, Primpeli A, Wichtlhuber B, Bizer C. Using schema.org annotations for training and maintaining product matchers. In: Proceedings of the 10th international conference on web intelligence, mining and semantics. WIMS 2020. Association for Computing Machinery, New York; 2020. p. 195–204. <https://doi.org/10.1145/3405962.3405964>.
- Peeters R, Bizer C. Dual-objective fine-tuning of bert for entity matching. *Proc VLDB Endow.* 2021;14(10):1913–21. <https://doi.org/10.14778/3467861.3467878>.
- Ristoski P, Petrovski P, Mika P, Paulheim H. A machine learning approach for product matching and categorization. *Semant Web.* 2018;9(5):707–28.
- Le Q, Mikolov T. Distributed representations of sentences and documents. In: Proceedings of the 31st international conference on machine learning-volume 32. ICML'14. JMLR.org; 2014. p. 1188–1196.
- Kozareva Z. Everyone likes shopping! multi-class product categorization for e-commerce. In: Proceedings of the conference of the North American chapter of the association for computational linguistics: human language technologies (NAACL). ACL; 2015. p. 1329–1333.
- Chavaltada C, Pasupa K, Hardoon D. A comparative study of machine learning techniques for automatic product categorisation. In: Proceedings of the international symposium on neural networks. Springer; 2017. p. 10–17.
- Xia Y, Levine A, Das P, Di Fabrizio G, Shinzato K, Datta A. Large-scale categorization of Japanese product titles using neural attention models. In: Proceedings of the 15th conference of the European chapter of the association for computational linguistics (EACL): volume 2, short papers. Association for Computational Linguistics; 2017. p. 663–668.
- Akritidis L, Fevgas A, Bozaris P. Effective unsupervised matching of product titles with k-combinations and permutations. In: IEEE

- 30th international conference on tools with artificial intelligence; 2018. p. 213–220.
30. Cevahir A, Murakami K. Large-scale multi-class and hierarchical product categorization for an e-commerce giant. In: Proceedings the 26th international conference on computational linguistics (COLING): technical papers. The COLING 2016 Organizing Committee; 2016. p. 525–535.
 31. Gupta V, Karnick H, Bansal A, Jhala P. Product classification in e-commerce using distributional semantics. In: Proceedings of COLING2016: technical papers. The COLING 2016 Organizing Committee; 2016. p. 536–546.
 32. Lee H, Yoon Y. Engineering doc2vec for automatic classification of product descriptions on o2o applications. *Electron Commer Res.* 2018;18(3):433–56.
 33. Borst J, Körner E, Opasjumruskit K, Niekler A. Language model cnn-driven similarity matching and classification for html-embedded product data. In: Proceedings of the semantic web challenge on mining the web of HTML-embedded product data co-located with the 19th international semantic web conference. CEUR Workshop Proceedings, vol. 2720. CEUR-WS.org; 2020.
 34. Li J, Dou Z, Zhu Y, Zuo X, Wen J-R. Deep cross-platform product matching in e-commerce. *Inf Retr J.* 2020;23:136–58.
 35. Zahera H, Sherif M. Probert: product data classification with fine-tuning bert model. In: Proceedings of the semantic web challenge on mining the web of HTML-embedded product data co-located with the 19th international semantic web conference. CEUR Workshop Proceedings, vol. 2720. CEUR-WS.org; 2020.
 36. Ha J-W, Pyo H, Kim J. Large-scale item categorization in e-commerce using multiple recurrent neural networks. In: Proceedings of the international conference on knowledge discovery and data mining (KDD). ACM; 2016. p. 107–115.
 37. Zhang Z, Paramita M. Product classification using microdata annotations. In: Ghidini C, Hartig O, Maleshkova M, Svátek V, Cruz I, Hogan A, Song J, Lefrançois M, Gandon F, editors. The semantic web-ISWC 2019. Cham: Springer; 2019. p. 716–32.
 38. Kim Y. Convolutional neural networks for sentence classification. In: Proceedings of the conference on empirical methods in natural language processing (EMNLP). Association for Computational Linguistics (ACL); 2014. p. 1746–1751.
 39. Yang L, E S, Xu S, Xiang Y. Bert with dynamic masked softmax and pseudo labeling for hierarchical product classification. In: Proceedings of the semantic web challenge on mining the web of HTML-embedded product data co-located with the 19th international semantic web conference. CEUR Workshop Proceedings, vol. 2720. CEUR-WS.org; 2020.
 40. Devlin J, Chang M-W, Lee K, Toutanova K. Bert: pre-training of deep bidirectional transformers for language understanding. 2018. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
 41. Kannan A, Givoni IE, Agrawal R, Fuxman A. Matching unstructured product offers to structured product specifications. In: Proceedings of the 17th ACM SIGKDD international conference on knowledge discovery and data mining. KDD '11. Association for Computing Machinery, New York; 2011. p. 404–412. <https://doi.org/10.1145/2020408.2020474>.
 42. Gopalakrishnan V, Iyengar SP, Madaan A, Rastogi R, Sengamedu S. Matching product titles using web-based enrichment. In: Proceedings of the 21st ACM international conference on information and knowledge management. CIKM '12. Association for Computing Machinery, New York; 2012. p. 605–614. <https://doi.org/10.1145/2396761.2396839>.
 43. Vandić D, van Dam J-W, Frasinćar F. Faceted product search powered by the semantic web. *Decis Support Syst.* 2012;53(3):425–37. <https://doi.org/10.1016/j.dss.2012.02.010>.
 44. van Bezu R, Borst S, Rijkse R, Verhagen J, Vandić D, Frasinćar F. Multi-component similarity method for web product duplicate detection. In: Proceedings of the 30th annual ACM symposium on applied computing. SAC '15. Association for Computing Machinery, New York; 2015. p. 761–768. <https://doi.org/10.1145/2695664.2695818>.
 45. Shah K, Kopru S, Ruvini J-D. Neural network based extreme classification and similarity models for product matching. In: Proceedings of the 2018 conference of the North American chapter of the association for computational linguistics: human language technologies, vol. 3 (Industry Papers). Association for Computational Linguistics, New Orleans-Louisiana; 2018. p. 8–15. <https://doi.org/10.18653/v1/N18-3002>.
 46. Tracz J, Wójcik PI, Jasinska-Kobus K, Belluzzo R, Mroczkowski R, Gawlik I. BERT-based similarity learning for product matching. In: Proceedings of workshop on natural language processing in e-commerce. Association for Computational Linguistics, Barcelona; 2020. p. 66–75.
 47. Petrovski P, Bryl V, Bizer C. Integrating product data from websites offering microdata markup. In: Proceedings of the 23rd international conference on world wide web. WWW '14 companion. Association for Computing Machinery, New York; 2014. p. 1299–1304. <https://doi.org/10.1145/2567948.2579704>.
 48. Petrovski P, Bizer C. Learning expressive linkage rules from sparse data. *Semant Web.* 2020;11(3):549–69.
 49. Köpcke H, Thor A, Thomas S, Rahm E. Tailoring entity resolution for matching product offers. In: Proceedings of the 15th international conference on extending database technology. EDBT '12. Association for Computing Machinery, New York; 2012. p. 545–550. <https://doi.org/10.1145/2247596.2247662>.
 50. Kiapour MH, Han X, Lazebnik S, Berg AC, Berg TL. Where to buy it: matching street clothing photos in online shops. In: 2015 IEEE international conference on computer vision (ICCV), 2015; p. 3343–3351. <https://doi.org/10.1109/ICCV.2015.382>.
 51. Wang X, Sun Z, Zhang W, Zhou Y, Jiang Y-G. Matching user photos to online products with robust deep features. In: Proceedings of the 2016 ACM on international conference on multimedia retrieval. ICMR '16. Association for Computing Machinery, New York; 2016. p. 7–14. <https://doi.org/10.1145/2911996.2912002>.
 52. Peeters R, Bizer C, Glavaš G. Intermediate training of BERT for product matching. In: Proceedings of the 2nd international workshop on challenges and experiences from data integration to knowledge graphs co-located with 46th international conference on very large data bases; 2020.
 53. Londhe N, Gopalakrishnan V, Zhang A, Ngo HQ, Srihari R. Matching titles with cross title web-search enrichment and community detection. *Proc VLDB Endow.* 2014;7(12):1167–78. <https://doi.org/10.14778/2732977.2732990>.
 54. Wu Y, Ngai EWT, Wu P, Wu C. Fake online reviews: literature review, synthesis, and directions for future research. *Decis Support Syst.* 2020;132:113280. <https://doi.org/10.1016/j.dss.2020.113280>.
 55. Mandhula T, Pabboju S, Gugulotu N. Predicting the customer's opinion on amazon products using selective memory architecture-based convolutional neural network. *J Supercomput.* 2020;75:5923–47.
 56. Gani A. Amazon sues 1,000 'fake reviewers' 2015. <https://www.theguardian.com/technology/2015/oct/18/amazon-sues-1000-fake-reviewers>. Accessed 12 Oct 2021.
 57. Salminen J, Kandpal C, Kamel AM, Jung S-G, Jansen BJ. Creating and detecting fake reviews of online products. *J Retail Consum Serv.* 2022;64:102771. <https://doi.org/10.1016/j.jretconser.2021.102771>.
 58. Xu Y, Yang Y, Han J, Wang E, Ming J, Xiong H. Slanderous user detection with modified recurrent neural networks in recommender system. *Inf Sci.* 2019;505:265–81.

59. Zhang L, Wu Z, Cao J. Detecting spammer groups from product reviews: a partially supervised learning model. *IEEE Access*. 2018;6:2559–68. <https://doi.org/10.1109/ACCESS.2017.2784370>.
60. Ball L, Elworthy J. Fake or real? the computational detection of online deceptive text. *J Market Anal*. 2014;2:187–201. <https://doi.org/10.1057/jma.2014.15>.
61. Ren Y, Zhang Y. Deceptive opinion spam detection using neural network. In: Proceedings of COLING 2016, the 26th international conference on computational linguistics: technical papers. The COLING 2016 Organizing Committee, Osaka; 2016. p. 140–150. <https://aclanthology.org/C16-1014>.
62. Yuan S, Wu X, Xiang Y. Task-specific word identification from short texts using a convolutional neural network. *Intell Data Anal*. 2018;22(3):533–50.
63. Liu Y, Pang B, Wang X. Opinion spam detection by incorporating multimodal embedded representation into a probabilistic review graph. *Neurocomputing*. 2019;366:276–83. <https://doi.org/10.1016/j.neucom.2019.08.013>.
64. Liu Y, Ott M, Goyal N, Du J, Joshi M, et al. RoBERTa: a robustly optimized BERT pretraining approach. 2019. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692).
65. Mikolov T, Sutskever I, Chen K, Corrado G, Dean J. Distributed representations of words and phrases and their compositionality. In: Proceedings of the 26th international conference on neural information processing systems-vol. 2. NIPS'13. Red Hook: Curran Associates Inc.; 2013. p. 3111–3119.
66. Beltagy I, Lo K, Cohan A. Scibert: a pretrained language model for scientific text. *arXiv preprint 2019*. [arXiv:1903.10676](https://arxiv.org/abs/1903.10676).
67. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, Kang J. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*. 2020;36(4):1234–40.
68. Klein G, Kim Y, Deng Y, Nguyen V, Senellart J, Rush AM. Opennmt: neural machine translation toolkit. 2018. [arXiv:1805.11462](https://arxiv.org/abs/1805.11462).
69. Mudgal S, Li H, Rekatsinas T, Doan A, Park Y, Krishnan G, Deep R, Arcaute E, Raghavendra V. Deep learning for entity matching: a design space exploration. In: Proceedings of the 2018 international conference on management of data. SIGMOD '18. Association for Computing Machinery, New York; 2018. p. 19–34. <https://doi.org/10.1145/3183713.3196926>.
70. Jiang N, de Marneffe M-C. Evaluating BERT for natural language inference: a case study on the CommitmentBank. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP). Association for Computational Linguistics, Hong Kong, China; 2019. p. 6086–6091. <https://doi.org/10.18653/v1/D19-1630>. <https://www.aclweb.org/anthology/D19-1630>.
71. Bojanowski P, Grave E, Joulin A, Mikolov T. Enriching word vectors with subword information. 2016. [arXiv preprint arXiv:1607.04606](https://arxiv.org/abs/1607.04606).
72. Potvin B, Villemaire R. Robust web data extraction based on unsupervised visual validation. In: Nguyen NT, Gaol FL, Hong T-P, Trawiński B, editors. Intelligent information and database systems. Cham: Springer; 2019. p. 77–89.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.