

Towards Open-World Product Attribute Mining: A Lightly-Supervised Approach

Liyan Xu^{1,2*} Chenwei Zhang³ Xian Li³ Jingbo Shang⁴ Jinho D. Choi²

¹Pattern Recognition Center, WeChat AI ²Emory University

³Amazon.com ⁴University of California, San Diego

liyan.xu@emory.edu {cwzhang,xianlee}@amazon.com

Abstract

We present a new task setting for attribute mining on e-commerce products, serving as a practical solution to extract open-world attributes without extensive human intervention. Our supervision comes from a high-quality seed attribute set bootstrapped from existing resources, and we aim to expand the attribute vocabulary of existing seed types, and also to discover any new attribute types automatically. A new dataset is created to support our setting, and our approach Amacer is proposed specifically to tackle the limited supervision. Especially, given that no direct supervision is available for those unseen new attributes, our novel formulation exploits self-supervised heuristic and unsupervised latent attributes, which attains implicit semantic signals as additional supervision by leveraging product context. Experiments suggest that our approach surpasses various baselines by 12 F1, expanding attributes of existing types significantly by up to 12 times, and discovering values from 39% new types. Our data and code can be found at <https://github.com/lxucs/wom>.

1 Introduction

Attribute mining (or product attribute extraction) is to extract values of various attribute types (e.g. *colors, flavors*) from e-commerce product description, which is a foundational piece for product understanding in online shopping services, enabling better search and recommendation experience.

Within this task regime, different settings have been studied. Most pioneer works deem it as a closed-world setting, where models are trained to identify a fixed set of pre-defined attribute types (Ghani et al., 2006; Putthividhya and Hu, 2011; Zheng et al., 2018), similar to the standard named entity recognition (NER). Recent works start to step up towards the open-world aspect that supports

*Primary work done as an applied scientist intern at Amazon, before joining WeChat AI.

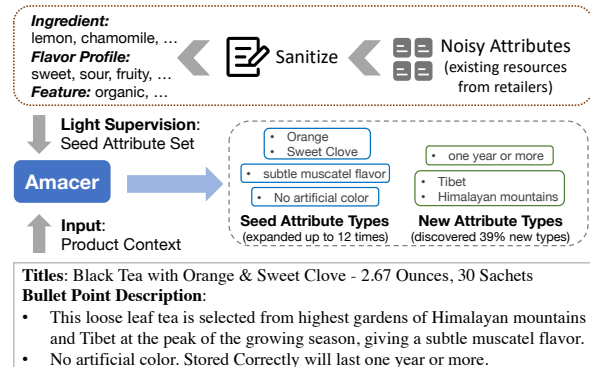


Figure 1: Illustration of our task setting on one product: given light supervision from seed attributes, our approach Amacer aims to expand attribute vocabulary of seed types, and to also discover values of any new types (*Shelf Life, Origin*) not covered by seeds. The outputs on all products are thus attribute clusters with diverse values. Evaluation is based on clustering metrics, as new clusters are not named beforehand.

extraction of new attribute types unseen in training. Particularly, several works have focused on the zero-shot perspective (Xu et al., 2019; Yang et al., 2022), enabling extraction of a new attribute type during inference if given a name or description of this new type, which is a more realistic setting to this task, as new types of products and attributes are constantly emerging in the real world.

In this work, we formulate the attribute mining task one step further towards the ultimate open-world setting: given product-related description, the objective is to identify as many new values of existing attribute types, as well as any new types that could be considered as reasonable attributes but not covered in training. As such, our setting automatically discovers new attributes, unlike the zero-shot setting that requires explicit specification of new types of interest. In addition, we also aim the model to work under limited supervision, by introducing only a relatively small seed attribute set in training, thereby remaining practical when only a few values are known for a certain attribute,

also for the fact that it is untenable to keep up high-coverage human annotations of ever-changing attributes, especially in e-commerce domain.

Figure 1 illustrates our overall task setting, where the model expands the attribute vocabulary of existing types, and discovers any new attributes, yielding numerous attribute clusters. A new dataset dubbed **WOAM (Weakly-supervised Open-world Attribute Mining)** is created to accommodate our setting, as described in Section 2. Targeting towards realistic open-world setting, our dataset covers full product horizons including titles and detailed description, where the latter provides rich context and is shown to contain more unseen attribute types than titles by 66% (Table 1). Moreover, distinguished from previous datasets that either require substantial annotation efforts (Zheng et al., 2018) or noisy distant-supervised data (Xu et al., 2019; Yang et al., 2022; Zhang et al., 2022), our training supervision comes from a high-quality seed attribute set constructed hybridly, combining data-driven and light human curation. Overall, our setting achieves good trade-offs with reasonable human interventions, under a practical scope with decent coverage on attributes.

We then propose our approach for this setting, dubbed **Amacer (Attribute mining with adaptive clustering and weak regularization)**. To overcome the challenge of limited supervision, we first introduce our approach to generate diverse spans of candidate attribute values from corpus (§3); then focus on representation learning by utilizing explicit supervision from seed attributes (§4), followed by the last step that performs grouping on candidate spans using refined, attribute-aware embeddings (§5). New formulations to mine more implicit semantic signals from product context are also proposed for new attribute discovery (§6).

Experiments on WOAM suggest that our approach outperforms various baselines by up to 12.5 F1. Furthermore, our novel formulation to leverage self-supervised and unsupervised semantic signals is shown effective to both existing and new attributes, especially boosting new attribute discovery by a good margin of 6.4 F1. Despite the limited amount of seed values, our model is able to expand the seed attribute vocabulary by up to 12 times (Table 15), and to discover values from 39% unseen attribute types on our test set. Overall, our contributions can be summarized as follows:

- We address a new setting in attribute mining as

a practical paradigm to extract open-world attributes under light human intervention.

- A new dataset is created, covering 66 attribute types with 42% unseen types from the seed set.
- A new approach is proposed to support our unique task setting, especially exploiting self-supervised and unsupervised semantic signals, which has not been explored by previous works.

2 Data

Our dataset WOAM consists of three parts, including: 1) text corpus; 2) seed attribute set for training; 3) human-annotated test set for evaluation. Full statistics of our dataset are provided in Table 11, and more details are provided in Appendix B.

Corpus Four common e-commerce product categories are included in our corpus: *Tea, Vitamin, Sofa, Phone Case*. For each category, we sampled 9,000+ products publicly listed on Amazon.com with full description available in English. Each product record can be represented as a tuple: (identifier, category, title, bullet points).

Seed Set For each category, the seed set consists of a few applicable attribute types (avg. 16.5 types per category) and their values (avg. 22 values per type). We adopt a hybrid approach for the construction: existing resources are first utilized to bootstrap the seed set, and human curation is performed upon to overcome the noisy issue existed in previous datasets (example shown in Table 10). Specifically, two steps are applied as below:

Automatic Sanitizing: we collect the raw product profiles that contain certain attributes provided by Amazon retailers, and perform frequency-based heuristics to heavily sanitize noisy attributes. First, long-tail attribute types that have fewer than 10 values are removed. Second, for each product category, if a unique value appears under multiple attributes types, we restrict it to only belong to its most common type. Lastly, for each attribute type, we only keep at most 100 values based on the top frequency, so to discard the tail values that we are less confident on. The resulting seed set thereby has a relatively small size but of higher quality after above three steps.

Human Curation: as the attribute set after sanitizing is relatively small, human curators can go through the entire set rather quickly and consolidate the final seed set (< 40 min per product category). Concretely, remaining noisy values are

	Type (New)	Value (New)	Tok	Gold
TT	46 (28%)	864 (70%)	20.1	5.7 (28.5%)
BP	65 (43%)	2787 (89%)	26.6	3.6 (13.8%)

Table 1: Characteristics of our dataset by titles (TT) and bullet points (BP) on the test set (full stats in Table 11): total number of unique attribute types/values, with the ratio of new types/values in parentheses; averaged number of tokens and gold values per title/bullet sequence, with the density of gold values per token in parentheses.

spotted and removed from their attribute types. Furthermore, granularity is adjusted such that ambiguous or coarse attribute types are split into multiple newly defined fine-grained types; similar attribute types are also merged into one type.

After we obtain the final seed set, we perform string match to obtain their occurrences in corpus, ready to be used for training. A development set is separately created that consists of sanitized profile attributes solely for hyperparameter tuning. Overall, our training supervision is built practically that balances between scalability and quality.

Test Set For each category, we collect additional products not covered in the raw corpus as the test set. Two in-house annotators are asked to annotate all spans that appear as reasonable attribute values of either an existing type from the seed set, or a brand-new type that fits the context. As with previous works, we do not allow overlapping spans: more complete spans are preferred over shorter and incomplete spans; each span is assigned a single attribute type that best describes its property.

Table 1 briefly specifies unique characteristics of our dataset. It is clear that most gold values are new values unseen from the seed set. Especially, bullet points have a higher ratio of new attribute types/values than titles, while those values are harder to extract due to longer text, sparser values, and more complex language structures. For comparison, our setting poses greater challenges than the most related previous dataset from a recent work OA-Mine (Zhang et al., 2022), which is under a much limited scope that consists of only titles with sparser and noisier seed attributes (detailed comparison is provided in Appendix A).

Our proposed approach for this dataset is presented in the following Section 3-6. Specifically, Section 3-5 introduce the overall pipeline depicted in Figure 2 that utilizes explicit signals from seed attributes, and Section 6 introduces our novel for-

mulation to exploit implicit signals beyond the limited seed attributes.

3 Candidate Span Generation

The first stage of our approach is to generate spans from product description that could be qualified as attribute values, producing a set of non-overlapping candidate spans, serving as a foundational step for this attribute extraction task.

With weak supervision in mind, this step should not simply rely on signals from the seed set; otherwise, it would become hard to generalize and lose diverse attribute expressions during inference. Therefore, directly employing a supervised model can be suboptimal. It is also tempting to use off-the-shelf phrase extraction tools such as AutoPhrase (Shang et al., 2018), however, the domain shift on e-commerce description of varied categories can severely affect recall, as observed by Zhang et al. (2022). The close work OA-Mine regards this stage as an unsupervised sentence segmentation task on product titles through language model probing (Wu et al., 2020), regarding each segment as a candidate span. Nonetheless, two shortcomings still remain. First, unlike titles, segmentation may not be suitable for bullet points, as most segments from bullet points would be noisy spans, demonstrated by the lower value density (13.8%) in Table 1. Second, being completely unsupervised, there is no task-specific adjustment in this process, suffering inadequate candidate quality.

In this work, we instead resort to a basic yet effective strategy that overcomes above issues, by using **syntax-oriented patterns**: we collect valid Part-of-Speech (POS) patterns for attribute values, and simply obtain all spans in the corpus that fit into those patterns as candidate spans, followed by rudimentary stopword filtering and overlapping span removal (prioritizing longer spans), yielding a smaller but higher-quality candidate set than that from sentence segmentation.

Valid POS patterns are acquired in a data-driven fashion without human intervention: we leverage the product profiles again, and obtain all POS sequences of their attribute values. These raw sequences are further compacted by removing consecutive duplicate POS tags, such that *healthy clean water* ([ADJ, ADJ, NOUN] → [ADJ, NOUN]) will share the same POS pattern as *clean water* ([ADJ, NOUN]). The resulting set of collected POS patterns serves to identify spans as well-formed or ill-formed phrases.

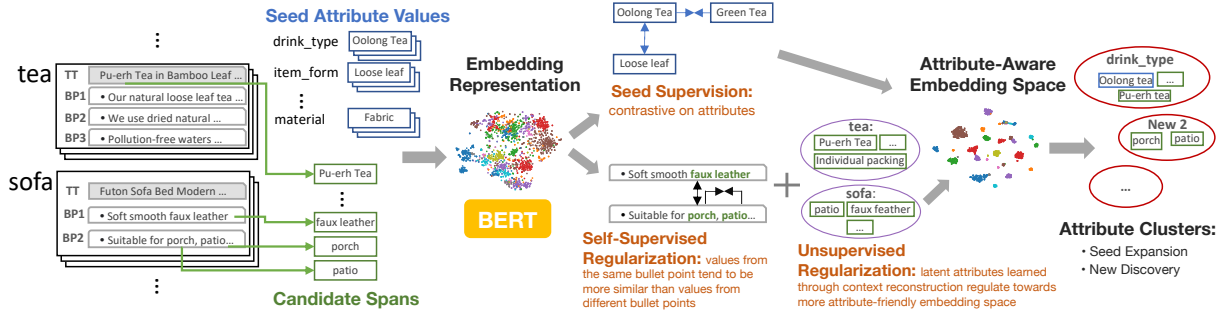


Figure 2: Illustration of our proposed approach Amacer. It generates candidate spans from product description (§3), and performs representation learning on embedding space, by utilizing: explicit supervision from seed attributes (§4); implicit semantic signals by self-supervised heuristic and unsupervised latent attributes (§6). Final attribute clusters can be obtained by grouping candidates through adaptive expansion and DBSCAN (§5).

Examples of our POS patterns are shown in Table 2. They regulate spans based on their syntactic features, without sole reliance on semantic supervision from the limited seed set, hence being able to capture diverse attribute expressions of vast variety. Overall, they serve as the quality guardrail for candidate spans, while reaping additional advantages: 1) easy to perform manual domain-specific adjustment; 2) scalable towards other product categories, as being data-driven; 3) efficient to run in practice.

<i>healthy clean water</i>	[ADJ, NOUN]	✓
<i>sweet and spicy taste</i>	[ADJ, CCONJ, ADJ, NOUN]	✓
<i>promotes healthy liver function</i>	[VERB, ADJ, NOUN]	✓
<i>are available during</i>	[VERB, ADJ, ADP]	✗
<i>freshness so every cup</i>	[NOUN, ADV, DET, NOUN]	✗

Table 2: Examples of POS patterns to recognize well-formed (✓) or ill-formed (✗) phrases.

As we depend on external tools to identify POS, this process is not without noises. Nonetheless, we find the empirical performance to be quite robust qualitatively. Moreover, it can be augmented with other techniques to mitigate noise in scenarios tailored to specific applications.

4 Explicit Signals for Seed Expansion

With both seed attribute values and candidate spans in-place, our next objective is to perform representation learning that refines the geometry of embedding space, such that values of similar attributes should have a closer embedding representation, and vice versa, as the key property to leverage in later grouping stage. In this section, we introduce the utilization of available seed attributes as **explicit supervision**, primarily targeting the vocabulary expansion of existing attribute types.

For each seed value or candidate span, we can have an initial representation on the embedding

space via encoding through pretrained language models such as BERT (Devlin et al., 2019). Concretely, we feed each text sequence (either a title or bullet point) to BERT, and obtain the contextualized representation of each span by averaging its token embedding, without introducing extra encoding parameters.

Supervised Contrastive Learning Contrastive learning is a natural fit to consume task signals from the seed set: for an anchor seed value v_a , a positive seed v_p from the same attribute, and a negative seed v_n from a different attribute, contrastive learning enforces (v_a, v_p) to be more similar than (v_a, v_n) on the embedding space. OA-Mine adopts a triplet loss (Schroff et al., 2015) for the supervised contrastive learning, as well as another regression loss (Reimers and Gurevych, 2019) that directly pushes the similarity of positive/negative pairs, requiring careful sampling and tuning. In our work, we simplify this supervised process by only using an in-batch negative contrastive loss (Khosla et al., 2020). Let I^s be all seed value indices, $P^s(i)$ be the indices of positive seeds that belong to the same attribute as seed i , $N^s(i) = I^s \setminus P^s(i)$ be the corresponding negative seeds. g_i is the L2-normalized embedding of seed i from the last layer of BERT encoding. The loss can then be denoted as:

$$\mathcal{L}^{su} = \sum_{i \in I^s} \frac{-1}{|P^s(i)|} \sum_{p \in P^s(i)} \log \frac{e^{(g_i \cdot g_p / \tau)}}{\sum_{j \in N^s(i)} e^{(g_i \cdot g_j / \tau)}}$$

τ is the temperature hyperparameter. As all embeddings are L2-normalized, $g_i \cdot g_j$ is effectively the cosine similarity as a distance measurement of two span representation. \mathcal{L}^{su} pushes seed values of the same attribute to have a similar representation, while pulling away seed values from different attribute types on the embedding space.

5 Candidate Span Grouping

After representation learning, a grouping stage upon candidate spans is followed. Each resulting cluster represents an attribute type, with each span inside being its attribute value. Unlike most related works that employ off-the-shelf clustering algorithms such as HAC, K-Means or DBSCAN (Elsahar et al., 2017; Zhao et al., 2021; Zhang et al., 2022), we propose a more fine-grained grouping strategy, which first explicitly addresses the expansion of existing seed attributes, then discovers new potential attributes, as described below.

Adaptive Expansion on Existing Attributes We borrow the concept from few-shot learning, and regard each existing seed attribute set as a support set. The distance between each candidate span c_i and each support set \mathcal{S}_j is measured by \mathcal{D} , which is the averaged cosine distance between the candidate and each seed values, as in Eq (1). A candidate c_i is added to an attribute j if $\mathcal{D}(c_i, \mathcal{S}_j) < \mathbf{t}_j$, where \mathbf{t}_j is a threshold calculated adaptively based on its support set, as in Eq (2). Particularly, $\delta \in (0, 1]$ is a hyperparameter to relax the threshold that can be tuned on the development set.

$$\mathcal{D}(c_i, \mathcal{S}_j) = \frac{1}{|\mathcal{S}_j|} \sum_{s_k \in \mathcal{S}_j} \text{cosine}(c_i, s_k) \quad (1)$$

$$\mathbf{t}_j = \delta \cdot \frac{1}{|\mathcal{S}_j|^2} \sum_{s_u, s_v \in \mathcal{S}_j} \text{cosine}(s_u, s_v) \quad (2)$$

More Attribute Coverage For remaining candidate spans, more clusters are mined to increase coverage primarily for potential new attributes. We also resort to off-the-shelf DBSCAN that can automatically discover clusters and distinguish noises based on the pairwise cosine distance.

The union of clusters from the above two stages serve as the final result of the candidate grouping.

6 Implicit Signals for New Discovery

Since the seed set only provides semantic signals regarding seed attributes, the majority of candidate spans lack proper supervision, as most of them are absent from the seed set, especially for those new attributes that have no direct supervision during representation learning. Therefore, it is desirable to exploit additional **implicit signals** towards more new-attribute-friendly embedding space, and we propose novel methods to tackle the challenge by fully leveraging product context through **self-supervised** and **unsupervised regularization**.

6.1 Self-Supervised Contrastive Learning

To utilize the product context, we formulate a self-supervised contrastive heuristic similar to skip-gram in word2vec (Mikolov et al., 2013). We regard each bullet point as a window: pushing two candidate spans within the same window (same bullet point) to have closer representation than two spans not in the same window (different bullet points of a product). It is based on the general observation that different bullet points usually discuss different product perspectives, but within each point, similar attributes or topics are usually mentioned. Though noisy, useful semantic signals could still be revealed given enough corpus, similar to the skip-gram training.

Let I^b be all candidate span indices in bullet points, $P^b(i)$ be the indices of positive spans within the same bullet point as i , $N^b(i)$ be the corresponding negative spans from different bullet points of the same product. The self-supervised contrastive loss is denoted as:

$$\mathcal{L}^{ss} = \sum_{i \in I^b} \frac{-1}{|P^b(i)|} \sum_{p \in P^b(i)} \log \frac{e^{(g_i \cdot g_p / \tau)}}{\sum_{j \in N^b(i)} e^{(g_i \cdot g_j / \tau)}}$$

We regard \mathcal{L}^{ss} as a form of regularization, assigning a small coefficient during training. The final loss is described as in Eq (8).

6.2 Unsupervised Latent Attributes

More useful signals could still be revealed from product context in addition to the bullet point heuristic. Inspired from topic modeling, e.g. Latent Dirichlet Allocation (LDA) (Blei et al., 2003), a classic generative method that discovers latent topics unsupervisedly from bag-of-words documents, here we propose a formulation of latent attributes to regulate the embedding space, providing implicit signals based on the semantic distribution of corpus, especially beneficial to new attribute discovery that has no direct supervision. We adapt the neural LDA work from Miao et al. (2017); Dieng et al. (2020), and regard topics as attributes in our setting. The main idea is that each product can be rendered as a composition of spans (equivalently, bag-of-spans) generated from different latent attributes based on the following two distributions.

Product-to-Attribute Distribution Given the context of a product, the model predicts a distribution over K latent attributes, where K is a hyperparameter. Latent attributes of higher probabilities

play a larger role in a product’s semantics. Since learning the true distribution is intractable, variational inference is applied such that we posit the distribution family to be multivariate Gaussian with diagonal covariance matrix, and fix the prior distribution as standard Gaussian (Dieng et al., 2020). Hence, the posterior Product-to-Attribute distribution can be obtained by simply predicting the mean and variance of multivariate Gaussian. Let p represent a product, \mathbf{h}^p be its context representation, μ_k^p/σ_k^p be its mean/variance for the latent attribute k predicted by the model. A sampled probability of attribute k for product p can be denoted as α_k^p :

$$\mu_k^p/\sigma_k^p = W_k^{\mu/\sigma} \cdot \mathbf{h}^p \quad (3)$$

$$\tilde{\alpha}_k^p \sim \mathcal{N}(\mu_k^p, \sigma_k^p) \quad (4)$$

$$\alpha_k^p = \text{softmax}(\tilde{\alpha}_k^p) |_{k=1}^K \quad (5)$$

$W_k^{\mu/\sigma}$ is a learned parameter to predict mean and variance. For \mathbf{h}_p , we use the averaged CLS representation of its product title and all bullet points.

Attribute-to-Span Distribution For each latent attribute, the model also learns a distribution over candidate spans; spans of high probabilities are the representatives of this attribute. Following Dieng et al. (2020), rather than building an explicit distribution, the model instead simply learns an attribute embedding, so that the distribution can be obtained by measuring the similarity of the attribute embedding and span embeddings. Let h_k be the k ’th attribute embedding learned by the model, g_c be the representation of a candidate span c , and \mathcal{C} be all unique candidate spans from all products in a training batch. The distribution of an attribute k over candidates \mathcal{C} can be denoted as:

$$\beta_{kc} = \text{softmax}(h_k \cdot g_c) |_{c \in \mathcal{C}} \quad (6)$$

Optimization Given the above two distributions for a product p , the model can easily get the **Product-to-Span** distribution $\mathcal{P}(c|p)$ by marginalizing out the latent attributes, as in Eq (7), which can then be used to optimize a reconstruction objective, such that spans actually appeared in product p should have higher probability than those who do not. Let $V(p)$ be the candidate spans in a product p , m be the total number of products. The unsupervised reconstruction loss \mathcal{L}^{un} can be estimated by

evidence lower bound (ELBO) as:

$$\mathcal{P}(c|p) = \sum_{k=1}^K \alpha_k^p \cdot \beta_{kc} \quad (7)$$

$$\mathcal{L}^{un} = - \sum_{p=1}^m \left(\sum_{c' \in V(p)} \log \mathcal{P}(c'|p) + \text{KL}(\tilde{\alpha}^p || \hat{\alpha}) \right)$$

$$\mathcal{L} = \mathcal{L}^{su} + \lambda^{ss} \cdot \mathcal{L}^{ss} + \lambda^{un} \cdot \mathcal{L}^{un} \quad (8)$$

where $\hat{\alpha}$ is the fixed standard Gaussian (prior Product-to-Attribute distribution). The first term of \mathcal{L}^{un} is the log-likelihood to encourage higher probability for actually appeared candidate spans in a product, and the second KL-divergence term regularizes the posterior attribute distribution $\tilde{\alpha}_p$ to be close to the standard Gaussian $\hat{\alpha}$.

The final loss \mathcal{L} during representation learning is constituted by three losses; λ^{ss} and λ^{un} are hyper-parameters that control the regularization strength.

7 Experiments

Experiments are conducted on our dataset in multiple model settings, including various baselines. Three different types of models are examined based on how attribute spans are obtained:

(1) Closed-world models based on sequence-tagging that extract spans upon predicted BIO tags of existing attributes, which do not support new attribute discovery natively. Two models are experimented: **Tx-CRF**, a generic Transformers-CRF tagging model; **SU-OpenTag** (Xu et al., 2019), a popular tagging-based attribute extraction model.

(2) Open-world models that rely on sentence segmentation to obtain candidate spans. We use the code released from OA-Mine to obtain all text segments for our dataset. Two settings are included: **OA-Mine** (Zhang et al., 2022); **Amacer***, a stripped version of our approach removing regularization and directly taking segments as candidates.

(3) Open-world models that employ our syntax-based candidate generation (§3). Five settings are included: **DBSCAN** that directly performs DBSCAN clustering without representation learning; **DBSCAN+AE** that adds our proposed adaptive expansion (§4); **OA-Mine*** that substitutes segmentation with our candidate spans; **Amacer**, our full proposed approach; and **Amacer^R** that only utilizes seed supervision without regularization in §6.

For candidate span generation, we use spaCy¹ to obtain POS tags; a total of 96 valid POS patterns

¹<https://spacy.io>

		Exact Match					Partial Match				
		Jaccard	ARI	NMI	Recall	F1	Jaccard	ARI	NMI	Recall	F1
Closed-World (Tagging)	Tx-CRF	92.5	95.4	95.8	20.0	32.8	78.2	85.3	86.7	30.5	44.2
	SU-OpenTag	70.1	78.8	87.1	22.1	34.5	61.7	72.6	79.5	34.7	46.6
Open-World (Segment)	OA-Mine	63.5	74.4	78.8	25.3	36.9	48.8	60.9	64.9	40.5	46.7
	Amacer*	69.9	78.0	84.1	29.0	41.7	58.4	68.8	73.7	47.8	54.9
Open-World (Syntax)	DBSCAN	22.4	29.8	69.5	17.3	23.6	20.6	24.7	60.7	26.9	30.3
	DBSCAN+AE	32.8	41.8	61.2	30.3	35.9	25.1	30.1	47.1	50.5	40.7
	OA-Mine*	55.8	68.2	73.6	30.8	41.1	40.6	52.0	57.2	50.1	49.8
	Amacer ^{-R}	58.3	69.6	79.2	35.5	46.3	46.3	57.6	65.8	57.7	56.9
	Amacer	67.2	76.9	84.0	35.7	47.6	52.7	63.8	70.4	57.1	59.1

Table 3: Evaluation results on the test set of our new dataset WOAM, with F1 being the overall evaluation metric. See Section 7 for detailed specifications of model settings and evaluation metrics. Each number is the macro-average across all product categories. Models with lower Recall tend to have higher Jaccard/ARI/NMI scores, as they produce fewer (and easier) attribute clusters of higher purity. The best performance by both Exact/Partial-F1 is the underlined score achieved by our approach Amacer (statistically significant from t-test > 95% confidence).

are acquired from product profiles (Section 3). The same BERT-Large is used as the encoder for all models. Our detailed hyperparameter settings are provided in Appendix C.

Evaluation Metrics Standard clustering evaluation metrics are used: **Jaccard**, Adjusted Rand Index (**ARI**), Normalized Mutual Information (**NMI**), to compare the attribute assignments on gold spans; **Recall**, to evaluate gold cluster coverage. As above metrics are consistent with OA-Mine, the evaluation adopts exact-match on predicted/gold spans. However, it could become over-restrictive as span boundaries can be quite subjective in this open-world setting, losing the information of near-correct predictions. Thus, we also provide a relaxed evaluation that allows partial-match on spans, such that a predicted span is considered an attribute value if more than half of the span falls into a gold value.

To assess the overall performance of a model, we roughly regard the averaged number of Jaccard, ARI and NMI as *pseudo precision*, and derive a single pseudo-**F1** score based on the clustering *precision* and *recall*, serving as the main evaluation metric of each approach.

Results Table 3 shows the evaluation results by all model settings. Our full proposed approach Amacer surpasses both SU-OpenTag and OA-Mine by a large margin (10+ Exact/Partial-F1), achieving the best performance on this task. Further observations and ablation study can be obtained as below.

• **Open-world models identify more attributes than closed-world models.** The two tagging-based models underperform OA-Mine-based models and

our Amacer-based models, with noticeably lower recall. It can be attributed to two factors. First, as all spans are obtained through tagging learned solely from the seed set, they lack the ability to accept more diverse attribute values not covered in training, not being able to generalize well under limited supervision. Second, new attributes are left untouched, unlike the open-world counterparts.

• **Adaptive expansion on seed attribute types is effective for candidate grouping.** By simply comparing DBSCAN with DBSCAN+AE, adaptive expansion is shown greatly improving the recall by 13-23% and overall performance by 10+%. On a side note, there is still a huge gap between DBSCAN+AE and Amacer, demonstrating the necessity to refine embedding space by representation learning.

• **Syntax-oriented generation obtains candidate spans of higher quality than segmentation.** Both OA-Mine* and Amacer^{-R} that apply syntax-oriented candidates outperform their segmentation-based counterparts OA-Mine and Amacer*, especially for exact-match that brings a gap of 4+ F1. Notably, our generation step takes under 10 minutes to process each category on CPUs, while the segmentation requires several hours on a GPU. Qualitatively, we found that the segmentation often over-divides sentences, yielding many noisy and incomplete phrases.

• **Seed supervision is more efficiently utilized by in-batch negative contrastive loss.** Compared to the triplet loss and regression loss adopted in OA-Mine*, the in-batch loss is not only simpler but also improves 5+ F1 in this task. We found the regression loss that pushes cosine similarity to 1/-1

	Seed / New	Title / BP	Gold
OA-Mine*	51.2 / 24.6	56.6 / 49.0	61.2
Amacer ^{-R}	64.5 / 39.8	61.2 / 56.9	69.8
Amacer	66.0 / 46.2	61.5 / 59.3	71.9

Table 4: Decomposed evaluation (Partial-F1) by: seed attribute types only (Seed) / new attribute types only (New); product titles only (Title) / bullet points only (BP). Gold shows the result by taking gold values directly as candidate spans. Full metrics are provided in Table 12-14 (Appendix D).

	Span (Exact)			Span (Partial)		
	P	R	F	P	R	F
OA-Mine*	31.0	38.3	34.2	52.8	64.8	58.1
Amacer ^{-R}	27.8	41.9	33.4	46.7	70.3	56.1
Amacer	33.5	40.5	36.4	54.9	65.5	59.3

Table 5: Evaluation of precision/recall/F1 (P/R/F) on the final extracted spans against gold values by exact/partial-match, regardless of the attribute types.

for pos/neg pairs can be too harsh for the embedding space, as certain attribute types are indeed more related and not completely independent.

- **Regularization (§6) is able to bring additional semantic signals** useful to shape the attribute-aware embedding space, as shown by the 2.2 Partial F1 improvement of Amacer upon Amacer^{-R}, where the unsupervised latent attribute formulation contributes around 70% improvement. We provide further quantitative and qualitative insights in Section 8-9.

8 Quantitative Analysis

To quantify the unique challenges of this task, we decompose the evaluation to examine two perspectives specifically:

- Performance on new attribute types (only open-world evaluation) compared to seed types (only closed-world evaluation).
- Performance on attribute values in bullet points compared to titles.

Table 4 shows that all models suffer performance degradation on new attribute types unseen in training, comparing with those existing seed types, which corroborates the expectation that open-world discovery remains a tough challenge owing to no direct supervision. It is noteworthy that our approach brings significant improvement on new attributes; especially, our proposed regularization in Amacer boosts performance on existing types by relatively 2.3% upon Amacer^{-R}, while the improvement on

new types is 16.1%, which fulfills our motivation to provide semantic supervision for those new attributes. Compared to OA-Mine*, our approach exhibits smaller relative gap between existing and new types, discovering 39% new types (Recall in Table 12).

For more traits of our corpus, all models struggle to keep up the performance on bullet points compared to titles, showing that they are indeed harder to extract from due to their characteristics (Table 1&9). Interestingly, our proposed regularization is also able to reduce the gap from 4.3 to 2.2 Partial-F1, which can be credited to both self-supervised heuristic and unsupervised latent attributes, as they both leverage the product context mainly from bullet points.

To detach the impact of candidate generation, we provide additional views to assess the representation learning and grouping performance. The last column of Table 4 shows evaluation by using gold values as candidate spans directly. It clearly strengthens the advantage of our proposed representation learning methods, as Amacer outperforms OA-Mine* by 10+ Partial-F1.

Table 5 further evaluates span extraction of predicted values against gold values. All models are shown quite low Exact-F1 scores (< 37) and low precision (< 34), leaving room for future improvement to extract more correct candidate spans under limited supervision.

9 Qualitative Analysis

Seed Attributes: our approach performs generally well on seed attribute types. Table 8 shows examples of discovered new values on a seed type *Flavor Profile* (also see Table 15). Amacer is able to extract sensible and diverse expressions, given only 6 seed values as supervision. Each proposed component makes evident contribution: the candidate generation can capture unseen long-tail spans, such as *floral with honey notes*, *delicate zesty*, while the representation learning and grouping together are effective recognizing similar attribute values. Nearly 80 new flavor values are identified on our test set, expanding its vocabulary by 12 times.

New Attributes: it is inevitably difficult to discover values of new types, as models possess little prior knowledge as regards. For error analysis, we found that for most of these new types, their values are either absent in the predictions, or grouped as other existing attributes mistakenly. Table 6 shows

Our **English Afternoon tea** combines Keemun tea from the **Anhui province** in **China** with **Ceylon tea** from **Sri Lanka**. **Keemun teas** are **smooth** and slightly **sweet in taste**, while **Ceylon teas** are **crisp and refreshing**.

Wormwood (**Artemisia absinthium**) is a **bitter herb** found in **Eurasia**, **North Africa**, and **North America**.

Table 6: Examples of extracted spans by Amacer on two bullet point description; the colors of spans represent three predicted seed attribute types: *drink type*, *ingredient*, *flavor profile*. According to the gold annotation, the following spans should belong to a separate attribute cluster (marked as a new non-seed attribute type “*region of origin*”): **Anhui province, China, Sri Lanka, Eurasia, North Africa, North America**. The model mistakenly predicts them as two existing attributes, showing that open-world attribute discovery remains a tough challenge to be solved under this task setting. On the other hand, it is still encouraging to see these spans being extracted and recognized as certain attributes, since the model has not seen any location-specific attributes directly from the seed set.

6 Selected Learned *Latent Attributes* by Each Column

living room	orange	oolong	no synthetic dyes	vitamin d3	moto g pure
navy love seats	purple clear	black tea	premium ingredients	kids vitamin c	12
tufted sofa	virtually invisible	green tea	artificial ingredients	vitamin b12	apple iphone
upholstered loveseat	brown hue	ti kuan yin oolong	vegetarian	amino acids	nokia x100
velvet sofa	warm neural	herbal tea	vegan and gluten free	folic acid	galaxy s21 fe

Table 7: Examples of several learned latent attributes, with top candidate spans from corpus at each column (high-probability spans in each Attribute-to-Span distribution). These learned latent attributes can represent certain concepts and provide additional semantic signals during representation learning, especially for new attributes.

<i>Flavor Profile</i>	
Seed (6)	Extracted (80+)
sweet	nutty
sweetened	floral with honey notes
unsweetened	earthy
sour	tangy and fruity
bitter	sweet and savory spice flavors
fruity	smokiness
	delicate zesty
	refreshingly tart herbal

Table 8: Sampled predictions on TEA products of the seed attribute *Flavor Profile* capturing diverse new values. Full examples are provided in Table 15.

an example of the latter case; however, it is still encouraging that these new values are extracted and recognized as certain attributes, rather than being neglected by the model, which partially achieves the open-world discovery objective.

Latent Attributes: Table 7 shows examples of learned latent attributes resulted by contrastive loss and topic modeling. They resemble certain “concepts” that regulate towards more attribute-friendly embedding space. However, we also observe that certain learned attributes are repetitive, such that their attribute embeddings have high cosine similarity. This behavior aligns with the previously discovered issue known as *topic collapsing* (Srivastava and Sutton, 2017), leading to deficient discovery. We do not particularly address it in this work, and

leave it for future research.

10 Conclusion

In this work, we present a new task setting as a practical solution to mine open-world attributes without extensive human intervention. A new dataset is created accordingly, and our proposed approach is designed for light supervision, especially by utilizing a high-quality seed set, as well as exploiting self-supervised and unsupervised semantic signals from the context. Empirical results show that our approach effectively improves discovery upon baselines on both existing and new attribute types.

11 Limitations

The scope of our approach is intended for our specific task setting, which is proposed as a practical solution to mine open-world attributes without heavy supervision, and has not been studied previously. Our approach does require an external dependency of a POS tagger, and assumes high POS tagging quality on English. Thankfully, there are POS tools publicly available with high performance, and are quite robust against domain shift, mostly fulfilling the assumption.

Our current candidate generation that utilizes syntax-oriented patterns does not check the semantics, which can be another limitation. It introduces noisy spans in the process, such as “*supports joint*

health & overall” (in Table 15). Future works could consider combining syntax with semantics to alleviate noisy spans.

References

- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(null):993–1022.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2020. [Topic modeling in embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 8:439–453.
- Hady Elsahar, Elena Demidova, Simon Gottschalk, Christophe Gravier, and Frederique Laforest. 2017. Unsupervised open relation extraction. In *The Semantic Web: ESWC 2017 Satellite Events*, pages 12–16, Cham. Springer International Publishing.
- Rayid Ghani, Katharina Probst, Yan Liu, Marko Krema, and Andrew Fano. 2006. [Text mining for product attribute extraction](#). *SIGKDD Explor. Newsl.*, 8(1):41–48.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 2410–2419. JMLR.org.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13*, page 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Duangmanee Putthivithya and Junling Hu. 2011. [Bootstrapped named entity recognition for product attribute extraction](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1557–1567, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. [Facenet: A unified embedding for face recognition and clustering](#). In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823.
- Jingbo Shang, Jialu Liu, Meng Jiang, Xiang Ren, Clare R. Voss, and Jiawei Han. 2018. [Automated phrase mining from massive text corpora](#). *IEEE Transactions on Knowledge and Data Engineering*, 30(10):1825–1837.
- Akash Srivastava and Charles Sutton. 2017. [Autoencoding variational inference for topic models](#). In *International Conference on Learning Representations*.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. [Perturbed masking: Parameter-free probing for analyzing and interpreting BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4166–4176, Online. Association for Computational Linguistics.
- Huimin Xu, Wenting Wang, Xin Mao, Xinyu Jiang, and Man Lan. 2019. [Scaling up open tagging from tens to thousands: Comprehension empowered attribute value extraction from product title](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5214–5223, Florence, Italy. Association for Computational Linguistics.
- Li Yang, Qifan Wang, Zac Yu, Anand Kulkarni, Sumit Sanghai, Bin Shu, Jon Elsas, and Bhargav Kanagal. 2022. [Mave: A product dataset for multi-source attribute value extraction](#). In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining, WSDM ’22*, page 1256–1265, New York, NY, USA. Association for Computing Machinery.
- Xinyang Zhang, Chenwei Zhang, Xian Li, Xin Luna Dong, Jingbo Shang, Christos Faloutsos, and Jiawei Han. 2022. [Oa-mine: Open-world attribute mining for e-commerce products with weak supervision](#). In *Proceedings of the ACM Web Conference 2022, WWW ’22*, page 3153–3161, New York, NY, USA. Association for Computing Machinery.
- Jun Zhao, Tao Gui, Qi Zhang, and Yaqian Zhou. 2021. [A relation-oriented clustering method for open relation extraction](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9707–9718, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. **Opentag: Open attribute value extraction from product profiles**. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, page 1049–1058, New York, NY, USA. Association for Computing Machinery.

A Previous Work

As the most related previous work to our proposed task setting is OA-Mine (Zhang et al., 2022), we found that their released dataset is not ideal nor practical to serve as the testbed for this setting, due to three drawbacks:

- The seed attribute set is too sparse: there are only five seed values provided for each attribute type, leading to insufficient attribute extraction and discovery.
- The seed attributes can be quite noisy; especially, certain values appear under multiple attribute types, presenting noise and ambiguity to the model training (example shown in Table 10).
- The corpus only consists of product titles, and lacks the full product description taxonomy such as bullet points, which can provide richer information regarding attributes and also require stronger inference capability. Detailed statistics of bullet point description compared to titles are provided in Table 9.

Our dataset explicitly addresses above issues, and is constructed to provide higher quality and richer context, as introduced in Section 2.

	Tok	Cand	Seed	Gold	Type (New)
TT	20.1	7.3	2.9	5.7	46 (28.3%)
BP	26.6	8.8	1.2	3.6	65 (43.1%)

Table 9: Statistics of our dataset WOAM that show more comparison between product titles (TT) and bullet point description (BP). Tok is the averaged number of tokens per sequence; Cand is the averaged number of generated candidates described in Section 3. Seed is the averaged occurrences of seed values per sequence, and Gold is the averaged occurrences of gold values in the test set. Type denotes the total number of attribute types in the test set, with parentheses indicating the ratio of new types that do not exist in the seed attribute set.

Seed Attribute Type	Seed Attribute Values
material feature	organic, gmo free, kosher, caffeine free, gluten free
specialty	organic, natural, herbal, caffeine free, kosher
special ingredients	organic, kosher, gluten free, matcha, cinnamon
diet type	gluten free, kosher, vegan, paleo, halal

Table 10: An example of seed attributes for TEA products from the dataset released by OA-Mine (Zhang et al., 2022). The provided seed attributes can be quite ambiguous, with many overlapping values in between. As this dataset is constructed in a distant-supervised way, the sub-optimal quality can hinder the model training to discriminate on different attributes. Our seed set adopts a hybrid approach combining data-driven and human curation, producing a practical and higher-quality attribute extraction.

B Dataset

Full statistics of our new dataset WOAM are provided in Table 11. Our dataset is publicly available under the Apache 2.0 License.

Corpus Our corpus consists of e-commerce product description from selected product categories, collected under permissions. We do not find concerns regarding privacy issues or discriminatory content.

Product Profiles In addition, we also document three detailed issues existed in product profiles that are addressed in our seed set construction: data sparsity, noisy attributes, coarse granularity. Thus, the raw profiles are unable to serve as the full supervision directly for this attribute extraction task.

- Our preliminary study shows that 80-90% human-identified attribute values are missing from the product profiles; along with the missing values, around 40% identified attribute types are also absent in the profiles, which aligns with the previous observations from Zhang et al. (2022). The sparsity of product profiles further cultivates our research motivation to enrich the product profiles by discovering new attributes automatically.
- Attribute values resided in profiles can be quite noisy, as there are no restrictions on what values that sellers could provide regarding their products. In extreme cases, many irrelevant values may be provided by sellers in efforts to boost their product search performance, which can disrupt the training and make the model insensible.
- Certain attributes may not be used directly due to their coarse granularity. For example, an attribute type STYLE can be too ambiguous for sellers such that it essentially becomes a superset of more fine-grained attribute values including colors, flavors, visual styles, materials, etc.

	Raw Text Corpus					Seed Attributes			Test Set Attributes		
	TRN	DEV	TST	BP	Toks	Types	Mdn/Avg	Occ	Types (New)	Values (New)	Occ
WOAM	209662	4647	1425	82.8%	25.5	36	9 / 27.0	1.5	66 (42.4%)	3382 (86.9%)	3.9
-TEA	49828	1094	524	82.0%	22.9	14	10 / 23.3	1.6	26 (46.2%)	1154 (86.3%)	3.7
-VIT	50298	1127	413	82.1%	24.1	15	25 / 37.4	1.7	22 (31.8%)	835 (81.2%)	3.5
-SOFA	55655	1228	240	83.8%	26.9	19	9 / 12.8	1.3	32 (40.6%)	775 (92.1%)	4.7
-CASE	53881	1198	248	83.3%	27.5	18	8 / 15.0	1.3	30 (40.0%)	703 (89.2%)	4.6

Table 11: Overall statistics of our created WOAM dataset, with breakdown of each product category: TEA, VIT (vitamin), SOFA, CASE (phone case). TRN/DEV/TST: number of text sequences (titles or bullet points) for the training/development/test set; BP: ratio of bullet point sequences; Toks: averaged number of tokens per sequence. For the seed set, Types: number of seed attribute types, with Mdn/Avg being the Median/Averaged number of values per type; Occ: averaged occurrences of seed values per sequence. For the test set, Types/Values: number of unique attribute types/values by human annotations, with parentheses indicating the ratio of new types/values unseen from the seed set; Occ: averaged occurrences of annotated values per sequence.

C Experimental Settings

For representation learning, BERT-Large (Devlin et al., 2019) is adopted as the encoder and we freeze all layers except for the last four layers, allowing for a larger batch size and faster training, which we found performs similar to finetuning the entire BERT. We use a batch size as 128, learning rate as 2×10^{-5} , linear-decay learning rate scheduler with warm-up ratio as 0.01, max gradient clipping norm as 1.

Other hyperparameters are searched on the development set; in our final Amacer model, we set the temperature $\tau = 0.1$ in the contrastive loss, and the number of latent attributes $K = 50$ (Section 6). In the final loss Eq (8), we set $\lambda^{ss} = 0.01$ and $\lambda^{un} = 0.02$, regarding them as weak regularization that mines additional semantic signals.

At the grouping stage, we set the relaxation $\delta = 0.8$ in adaptive expansion Eq (2). For DBSCAN, we use the implementation from sklearn², and set eps as 0.05, min_samples as 4.

All training is conducted on a Nvidia Tesla V100 GPU with 32GB memory, and takes around 1 hour to finish each model.

D Quantitative Analysis

Full evaluation metrics are provided in Table 12 and 13, in regard to the quantitative analysis in Section 8. In particular, Table 12 separately shows the detailed evaluation results on existing seed attribute types only or on new attribute types only. Table 13 separately shows the detailed evaluation results on product titles only, or on bullet point

description only.

Table 14 shows the full evaluation metrics when using gold spans as candidate spans directly. Since all resulting spans will be gold values, the evaluation scores are the same for either partial-match or exact-match.

²<https://scikit-learn.org/stable/modules/generated/sklearn.cluster.DBSCAN.html>

		Exact Match					Partial Match				
		Jaccard	ARI	NMI	Recall	F1	Jaccard	ARI	NMI	Recall	F1
<i>Seed</i>	OA-Mine*	50.5	64.8	73.8	33.1	43.4	36.6	50.8	61.5	52.8	51.2
	Amacer ^{-R}	70.7	81.6	86.3	40.8	53.9	52.3	66.8	74.8	64.4	64.5
	Amacer	73.4	83.5	88.1	41.3	54.8	55.8	69.9	77.3	64.5	66.0
<i>New</i>	OA-Mine*	13.9	16.3	52.5	17.7	21.6	11.4	12.7	46.9	25.6	24.6
	Amacer ^{-R}	15.7	20.2	63.3	29.1	31.0	15.6	19.2	61.6	48.8	38.8
	Amacer	37.9	49.1	77.5	27.0	36.2	40.2	52.9	75.3	39.3	46.2

Table 12: Decomposed evaluation results on seed attribute types only (*Seed*) or on new attribute types only (*New*). All models have performance degradation on new attribute types, showing that discovering open-world new attributes is a harder task than extracting seed attribute types seen in the training.

		Exact Match					Partial Match				
		Jaccard	ARI	NMI	Recall	F1	Jaccard	ARI	NMI	Recall	F1
<i>Title</i>	OA-Mine*	78.0	84.8	85.4	35.6	48.4	44.8	55.1	61.8	59.6	56.6
	Amacer ^{-R}	84.0	89.1	88.2	37.6	50.9	54.1	64.1	68.4	60.8	61.2
	Amacer	84.3	88.9	90.3	38.0	51.5	56.1	65.5	70.0	60.0	61.5
<i>BP</i>	OA-Mine*	52.8	65.0	71.8	29.6	39.6	41.9	53.2	58.6	47.7	49.0
	Amacer ^{-R}	55.2	66.4	77.2	35.0	45.3	46.3	57.7	66.5	57.2	56.9
	Amacer	65.3	75.1	82.7	35.2	46.8	54.3	65.4	71.9	56.3	59.3

Table 13: Decomposed evaluation results on product titles only (*Title*) or bullet point description only (*BP*). All models show performance degradation on bullet point description, indicating that bullet point description has its own traits compared to titles, requiring stronger span extraction and inference.

		Jaccard	ARI	NMI	Recall	F1
<i>GOLD</i>	OA-Mine*	54.8	66.7	68.7	59.7	61.2
	Amacer ^{-R}	63.1	73.8	78.5	68.4	69.8
	Amacer	70.0	78.9	83.0	68.2	71.9

Table 14: Evaluation results by directly using gold attribute values as candidate spans. The overall evaluation of each model gets boosted as expected, and directly reflects the performance of our proposed representation learning and grouping (Section 4&6). Note that scores are the same for either exact-match or partial-match.

<i>Flavor Profile</i>		<i>Health Benefit</i>	
Seed (6)	Extracted	Seed (34)	Extracted
sweet	nutty	anti aging	supports nervous system health
sweetened	bold	anti-aging	hypoallergenic
unsweetened	savory tea	boost energy	relieve fatigue
sour	refreshing taste	cleansing	curb sugar and hunger cravings
bitter	warm	cold relief	thirst quenching bottle
fruity	great tasting beverage	detox	help support a healthy heart
	delicious drink	detoxification	brighten our day
	fruit-flavored	detoxify	helps boost metabolism
	floral with honey notes	digestive health	consistency
	flowery	energizer	lives
	earthy	fertility	environment
	tangy and fruity	gut health	promotes healthy liver function
	tart	head relief	nourishes
	delicate flavor	hydrated	awakening
	light	immune support	relieves gas and bloating
	slightly sweet and spicy	immunity	supports the cardiovascular system
	minty	laxative	supports joint health & overall
	hot or cold	metabolism	experiencing the true taste
	savored	moisturize	celebration
	sweet and savory spice flavors	mood tonic	hormone balance and reproductive health
	spicy taste	nausea relief	taste and active properties
	tasting	night cleanse	helps regulate female hormone function
	unsweetened zero calories	nourishing	relieve menopause symptoms
	unsweetened green tea flavor	reduce bloating	enhance libido
	sweet in taste	relaxing herbal	reduce pain
	crisp and refreshing	sinus soother	increase fertility
	smooth	sleep support	improve mood
	sweet and spicy taste	slenderizer	clear your head
	vegetal flavor	soothing	yet soothing
	smokiness	stress relief	helps support a healthy lifestyle
	bright and floral flavor	supports immune	energy & immunity booster
	complex and rich flavors	throat tamer	properties and ayurvedic benefits
	aromatic	weight loss	exceptional nutritious properties
	rich treat	weight management	nutritious
	plain		promote healthier lifestyle choices
	rich flavour		unique energy characteristics
	teas-malty		reduce the jitters and crash
	hearty		steady and prolonged alertness
	rich flavor		boost cognitive function
	toasty texture		body breakthrough trim
	delicately floral		balanced diet plan
	fruity flavor		improves eye and vision health
	slightly tangy		five senses
	delicate zesty		creating health & wellness foods
	accented		support healthy menstruation
	refreshingly tart herbal		release harmful toxins
	vibrant		morning cleanse
	pleasantly roasted		cleanse your digestive tract
	bitter notes		detoxify your whole body

Table 15: Sampled predictions on TEA products of two seed attributes: *Flavor Profile*, *Health Benefit*. **Seed** columns display all seed values of the two attributes; **Extracted** columns show the predictions, which are extracted spans by Amacer from product titles or bullet point description. Given the limited amount of seed values, the model is able to expand much more diverse and long-tail expressions of attributes of interest, by up to 12 times for *Flavor Profile* on the test set. Indeed, the predictions also contain certain noise, due to the lightly-supervised setting.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 10.
- A2. Did you discuss any potential risks of your work?
Our work is only intended for our task scope that extracts attributes from product description, and only requires publicly available e-commerce data without targeting specific groups. Our work does not require large computational resources neither. Therefore, we don't see risks directly related to our work.
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract, Section 1.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 2, Section 6, Appendix C.

- B1. Did you cite the creators of artifacts you used?
Section 6
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Appendix B. We will also add specific copyright and terms upon dataset release.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Appendix B. Our dataset will be publicly available.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Appendix B.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We covered the domains and languages in Section 2. Our dataset does not involve specific demographic groups.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 2, Appendix B.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C Did you run computational experiments?

Section 6.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Appendix C.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix C.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 6.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Appendix C.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Section 2.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

We will publish these details upon dataset release.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Our annotators are in-house researchers involved in this work.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

We use in-house data collected under permissions.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

We have obtained approval from an ethics review board.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Our dataset consists of e-commerce products and does not have direct concerns regarding specific groups/people.