

Consistent Text Categorization using Data Augmentation in e-Commerce

Guy Horowitz^{1,*}, Stav Yanovsky Daye², Noa Avigdor-Elgrabli², and Ariel Raviv²

¹Technion – Israel Institute of Technology

²Yahoo Research

guy.h@campus.technion.ac.il

{stav.yanovsky,noaa,arielr}@yahooinc.com

Abstract

The categorization of massive e-Commerce data is a crucial, well-studied task, which is prevalent in industrial settings. In this work, we aim to improve an existing product categorization model that is already in use by a major web company, serving multiple applications. At its core, the product categorization model is a text classification model that takes a product title as an input and outputs the most suitable category out of thousands of available candidates. Upon a closer inspection, we found inconsistencies in the labeling of similar items. For example, minor modifications of the product title pertaining to colors or measurements majorly impacted the model’s output. This phenomenon can negatively affect downstream recommendation or search applications, leading to a sub-optimal user experience.

To address this issue, we propose a new framework for consistent text categorization. Our goal is to improve the model’s consistency while maintaining its production-level performance. We use a semi-supervised approach for data augmentation and presents two different methods for utilizing unlabeled samples. One method relies directly on existing catalogs, while the other uses a generative model. We compare the pros and cons of each approach and present our experimental results.

1 Introduction

In the last two decades, widespread use of e-commerce platforms such as Amazon and eBay has contributed to a substantial growth in online retail. Such platforms rely on both explicit and implicit product features in order to deliver a satisfying user experience. There, the inferred product category is typically a crucial signal for many application such as browsing, search and recommendation.

We focus on improving an existing product categorization model, we refer to as ‘the categorizer’,

*The work was carried out during an internship at Yahoo Research.

that is employed by our company for fast categorization of billions of items on a daily basis. It classifies e-commerce items, such as products or deals, based on a predefined hierarchy of categories, namely GPT (Google Product Taxonomy). Given a product title, the categorizer assigns the most appropriate label in the taxonomy. The model itself is highly scalable and effective, so it is well-suited for settings with large and rapidly growing item catalogs. In our company, the categorizer is used as a standalone component in various e-commerce related services, such as recommendation, search, and ad ranking.

A recent examination of the categorizer’s output revealed inconsistencies in the labeling of similar items. It was evident that in some cases small variations in product titles, such as those relating to colors or measurements, significantly affect the categorizer’s output. This inconsistency negatively impacts search and recommendation algorithms that rely on the inferred category, leading to a poor user experience.

The concept of consistency in NLP tasks has been studied in various research works, including robustness to paraphrasing (Elazar et al., 2021) and robustness to adversarial attacks (Jin et al., 2020; Wang et al., 2020). Other works relate consistency issues with the misuse of spurious features during the learning phase (Arjovsky et al., 2019; Veitch et al., 2021; Wang et al., 2021).

When examining the performance of the categorizer in terms of accuracy alone, the inconsistency issue may be overlooked. But, since many recommendation pipelines depend on the output of the product categorizer, an inconsistent model can have severe implications on the user experience. In most cases, the differences include returning the parent category or a sibling category, rather than a completely different category path.

To tackle this inconsistency problem, we use different *data augmentation* techniques and enrich

the training data with item versioning, leading to a more consistent model. Data augmentation for improving various NLP tasks has been widely studied and surveyed (Shorten et al., 2021), and particularly in the context of consistency (Xie et al., 2020). Generating such data, both manually (Kaushik et al., 2019) and automatically (Rizos et al., 2019; Bari et al., 2020; Kumar et al., 2020), has shown to contribute to the robustness of learnt models in different settings. We chose to use data augmentation, without changing the current architecture of the already-in-use product categorizer for two main reasons. First, for scalability reasons, any change in the architecture might degrade the model’s ability to infer the categories of billions of items per day. Second, maintaining the current model architecture expedites the productization process and requires only minimal engineering effort.

This work defines a new framework, *Consistent Semi-Supervised Learning (Consistent-SSL)*, for consistent text categorization in the context of e-commerce (Section 2). We use an unlabeled clustered dataset as a source of legit item versioning. The dataset is derived from product catalogs, and includes clusters of different versions of items. We present two different methods to utilize this unlabeled clustered data: a self-training method and a generative approach (Section 3). We describe the datasets and the experimental framework we use for the evaluation of the proposed methods (Section 4). Finally, we detail results, showing an improvement in the consistency rate of 4-10% above the baseline model, and discuss the advantages and weaknesses of each method (Section 5).

2 Consistent Semi-Supervised Learning

We now formalize our notion of consistent classification and introduce the settings for consistent Semi-Supervised Learning (consistent-SSL).

2.1 Consistent Classification

In order to formalize consistent classification, let \mathcal{X} be our set of items, and $\mathcal{Y} = [c]$ for $c \in \mathbb{N}$, be a final set of labels. Each item $x \in \mathcal{X}$ corresponds to a label $y \in \mathcal{Y}$.

Additionally, let $\mathcal{V} : \mathcal{X} \rightarrow \mathcal{X}$, be a non-deterministic perturbation function which transforms an item from one version x to another \hat{x} . For example, if $x = \text{"blue T-shirt small size"}$, $\hat{x} \sim \mathcal{V}(x)$ could be $\hat{x} = \text{"black T-shirt small size"}$ or $\hat{x} = \text{"blue T-shirt large size"}$. We assume that

the perturbation function is label-preserving, i.e. $x, \hat{x} \sim \mathcal{V}(x)$ share the same label y . Let $p(x, y)$ be a joint distribution over items and labels and $p(x)$ the marginal distribution over items. The goal of consistent classification is to learn a classifier $f : \mathcal{X} \rightarrow \mathcal{Y}$ from a class F with a dual objective: a high expected **accuracy**, i.e. high expected value of the indicator that an item $x \in X$ is labeled by f to its correct label y :

$$\mathbb{E}_{(x,y) \sim p(x,y)} [\mathbb{1} \{f(x) = y\}] \quad (1)$$

and a high expected **consistency**, which we define as:

$$\mathbb{E}_{\substack{x \sim p(x), \\ \hat{x} \sim \mathcal{V}(x)}} [\mathbb{1} \{f(x) = f(\hat{x})\}] \quad (2)$$

i.e. the expected value of the indicator of two items $x, \hat{x} \sim \mathcal{V}(x)$ to be transformed by f to the same label. Therefore, the dual objective of f can be formalized as:

$$\min_f \mathbb{E}_{\substack{(x,y) \sim p(x,y), \\ \hat{x} \sim \mathcal{V}(x)}}} [\mathbb{1} \{f(x) \neq y\} + \lambda \mathbb{1} \{f(x) \neq f(\hat{x})\}] \quad (3)$$

where $\lambda \in \mathbb{R}$ controlling the balance between the accuracy loss and the consistency loss.

Note that there could be a trad-off between the accuracy objective and the consistency one. A model that is trained to disregard specific features like color or size would be more consistent but as those features might be informative in order to partition between some categories this could harm the overall accuracy. For example, if the color purple is more likely to appear in sport shoes than in evening shoes, a model that is trained to give less weight to colors may have a harder time distinguishing between sports and evening shoes while being more robust to changes in colors and thus more consistent.

2.2 Consistent-SSL Settings

In SSL settings, we are given labeled data $\mathcal{D}_L = \{(x_i, y_i)\}_{i=1}^l$, which is assumed to be sampled i.i.d. from p , and unlabeled data $\mathcal{D}_U = \{x_i\}_{i=l+1}^{l+u}$ possibly sampled from another distribution q . We tune a classifier f using both \mathcal{D}_L and \mathcal{D}_U .

This work extends the standard SSL settings to consistent-SSL. The unlabeled data \mathcal{D}_U is clustered with respect to the perturbation function \mathcal{V} , i.e. it consists of u sets of items X_i , each set contains k_i versions $\hat{x}_j^{(i)} \sim \mathcal{V}(x_i)$ of the same item x_i .

More formally, $\mathcal{D}_U = \{X_i\}_{i=l+1}^{l+u}$, where, $X_i = \{\hat{x}_j^{(i)}\}_{j=1}^{k_i}$, and $\hat{x}_j^{(i)} \sim \mathcal{V}(x_i)$ for $j = 1 \dots k_i$.

The goal in consistent-SSL is to learn a classifier f that optimizes the objective in Eq. (3) given \mathcal{D}_L and \mathcal{D}_U . Note that \mathcal{V} is unknown, and only appears indirectly in the \mathcal{D}_U samples.

3 Methods

We present two methods for consistent-SSL, Consistent Self Training (CST) and Consistent Generative Augmentation (CGA). Both methods utilize the unlabeled samples from \mathcal{D}_U for data augmentation. In each method we create an augmented set \mathcal{D}_{aug} using \mathcal{D}_U and train a classifier f on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. This approach optimizes indirectly the objective of Eq. (3), as we add additional training samples \mathcal{D}_{aug} that consists of different versions of the same items. The goal is to expose f to a more diverse set of item versions in training time, making it more robust to minor changes.

Let us review our approach using an illustrative example. Consider a dataset that contains clothing items. Assuming that \mathcal{D}_L , which was sampled from the distribution p , exhibits a spurious correlation between color of an item to its category (e.g. most of the black items are coats and most of the red items are dresses), then a classifier that was trained solely on \mathcal{D}_L will tend to rely on the color of the item when it predicts its category. When applying the model, \mathcal{V} could change the items' colors and therefore the classifier will not be consistent (e.g. if \mathcal{V} transforms a black coat to a red one, the classifier might predict different categories). But, assuming the training data includes an item in multiple colors (e.g. black coat, red coat, blue coat, etc.), with the same label (e.g. *Coats & Jackets*), then a model that is trained on such data will not relate a specific color to a specific label. Such a model will be encouraged to ignore the color of an item when it predicts the label, and therefore will be more robust to changes in color. Note that colors here are only an example of one kind of versioning of items. Spurious features in the data could be related to colors, measurements, models, materials etc.

3.1 Consistent Self Training (CST)

In our first method, named Consistent Self Training (CST), we add samples from \mathcal{D}_U to the labeled training data \mathcal{D}_L and a new classifier f is trained on the unified dataset. Since the data of \mathcal{D}_U is unlabeled, we perform a variant of self training




Figure 1: Illustration of CST pipeline. A base model f^{base} is trained on the labeled training set \mathcal{D}_L . It is then used to assign pseudo labels for the unlabeled samples from \mathcal{D}_U to create \mathcal{D}_{aug} . A classifier f is trained on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$.

(Lee et al., 2013; Arazo et al., 2020; Triguero et al., 2015). To make sure that \mathcal{D}_{aug} is consistent, it's important that each item set X_i is assigned with the same pseudo-label \tilde{y}_i . To calculate \tilde{y}_i , we first train a base model f^{base} on the labeled data \mathcal{D}_L and then use it to choose a single pseudo-label for each example set X_i , i.e. $\tilde{y}_i \leftarrow h(X_i; f^{\text{base}})$, where h is a function that given a set of examples and a classifier f^{base} returns a single label. For example, h could return the prediction of f^{base} that got the highest confidence score, or the most frequent prediction across X_i . The function h is an hyper-parameter of the method. Finally, a classifier f is trained over $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. Figure 1 shows an illustration of the CST pipeline, and a full description of the algorithm is presented in Appendix A.

3.2 Consistent Generative Augmentation (CGA)

We now detail our second method, we refer to as Consistent Generative Augmentation (CGA). Here, we train a generative model \mathcal{M} on \mathcal{D}_U in order to learn the perturbation function \mathcal{V} , and we use it to generate new samples based on the instances of \mathcal{D}_L . For this end, an item-pair dataset of different versions of items, $\mathcal{D}_{\text{pairs}}$ is constructed from \mathcal{D}_U ; $\mathcal{D}_{\text{pairs}} = \left\{ \left(\hat{x}_j^{(i)}, \hat{x}_{j'}^{(i)} \right) \mid l+1 \leq i \leq l+u \wedge j, j' \in [k_i] \right\}$. We train \mathcal{M} on $\mathcal{D}_{\text{pairs}}$ to generate the second item given the first of each pair, while maintaining its label. Note that $\hat{x}_{j'}^{(i)} \sim \mathcal{V}(\hat{x}_j^{(i)})$. Next, we generate an augmentation set \mathcal{D}_{aug} using \mathcal{D}_L by applying \mathcal{M} on each $(x, y) \in \mathcal{D}_L$ to get a new labeled sample (\hat{x}, y) . Note that we can use \mathcal{M} to generate multiple new samples from a single sample x . After creating \mathcal{D}_{aug} , we filter it using a score function $s : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ that aims to measure the quality of the generated \hat{x} with respect to its origin x . Additionally, we remove low quality samples from \mathcal{D}_{aug} according to some predefined filter threshold T . Finally, we train a classifier f over $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. Both s and T are hyper-parameter

of the CGA method. Figure 2 shows an illustration of the CGA pipeline, and a full description of the algorithm is presented in Appendix B.




Figure 2: Illustration of CGA pipeline. A generative model \mathcal{M} is trained on pairs of items from the catalog dataset \mathcal{D}_U . Then it is used to augment the labeled training set \mathcal{D}_L . The generated samples are filtered using a score function s . A classifier f is trained on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$.

3.3 Methods Comparison

We compare the two proposed methods by three main aspects: the quality of the augmented product titles, the quality of the labels and the overall distribution.

Considering the quality of the product titles, the CST method utilizes the unlabeled clustered data itself and thus provides product titles that are sampled from the real world and captures information about the true perturbation function \mathcal{V} . In contrast, the CGA method uses generated product titles, which may not represent \mathcal{V} accurately. Regarding the label quality, the CGA method utilizes labels that are taken directly from the ground truth labels of the original items and thus of a better quality than the ones of the CST method, which uses calculated "pseudo-labels". With respect to the distribution of the data, the generated samples in the CGA method are taken directly from the distribution p of the labeled training set. In contrast, in CST the unlabeled data comes from a distribution q that is different than p , thus biasing the overall distribution of the training set.

The quality of the product titles in the augmentation set impacts the consistency and corollary the overall optimization of the model f . On the other hand, both the quality of the labels and the distribution of the augmentation set influence the accuracy which again affects the overall optimization of f .

4 Empirical Evaluation

We now present our experimental results. We note that in all of our experiments, we use a model that is based on FastText (Joulin et al., 2016) architecture, and has an hierarchical structure. This specific model is found to perform well on our task, as it

takes into account the hierarchical structure nature of the labels. For more details, see Appendix C.

4.1 Train And Test Data

We conduct experiments using an e-commerce text classification dataset in order to empirically evaluate our methods. The items in this dataset are titles of commercial products, represented as free text, and the labels are the items' categories. The labels are taken from a hierarchic products taxonomy with 4 levels of granularity $\{L_i\}_{i=1}^4$. For example, consider a product title such as "*Greenies Breath Buster Bites Fresh Flavor Grain-Free Dental Dog Treats, 1.2-oz bag*", and its corresponding category *Animals & Pet Supplies > Pet Supplies > Dog Supplies > Dog Treats*.

Our dataset contains 184k labeled samples with 3k different labels, and additional 1.3M unlabeled samples. The labeled samples correspond to real-world commerce related items, and are labeled by human annotators. The unlabeled samples are retrieved from a product catalog of multiple retailers that includes grouping information. Each group contains multiple versions of the same item, e.g. "*L.A. Girl, Matte Lipstick, Snuggle, 0.10 oz*" and "*L.A. Girl, Matte Lipstick, Bite Me, 0.10 oz*". There are 363k different groups in the unlabeled catalog data, each group contains 2 to 192 items, and the average group size is 3.6. We note that the labeled and unlabeled data sets originate from different sources. This results in different category representation between the labeled and unlabeled data, e.g. several categories in the unlabeled data have low coverage compared to the labeled one.

Our experiments measure both accuracy and consistency of the tested models. To this end, we create two different test sets:

Accuracy test. The accuracy test is a standard test set that consists of labeled samples, on which we compute the weighted average F1 score of a given model. The accuracy test contains 23k labeled examples sampled uniformly at random from the labeled data. We use the remaining 161k labeled samples as the \mathcal{D}_L .

Consistency test. The consistency test consist of pairs of item titles (\hat{x}^1, \hat{x}^2) , each pair includes two different versions of the same item. We define the *consistency rate* of a given model f to be the percentage of the (\hat{x}^1, \hat{x}^2) pairs from the consistency test that receive the same label prediction by f , i.e. $f(\hat{x}^1) = f(\hat{x}^2)$. We create this test set by




Figure 3: Distributions of the different versions of the data for CST. The labels are presented in L_1 granularity.




Figure 4: CGA experiments results.

sampling 9k groups from the unlabeled data, then by sampling one pair of different titles (\hat{x}^1, \hat{x}^2) from each group. Since the consistency rate of a model on this test should be an empirical evaluation of its consistency as defined in Eq. (2), the distribution of the data in this test should be similar to the distribution of the data in the accuracy test. To mitigate some of the discrepancy between the unlabeled and labeled datasets, we sub-sample the unlabeled dataset according to the L_1 distribution of the labeled set. We use the unlabeled samples that are not selected for the consistency test as \mathcal{D}_U for training.

4.2 Experimental Framework

This subsection describes in detail the configuration of the proposed methods, and the baselines that were used for comparison.

4.2.1 Baselines

For the first **Baseline** model, we use the existing product categorization model, trained using only \mathcal{D}_L . The second baseline is a **ColorsSizes-Blind (CS-Blind)** model. We train it using \mathcal{D}_L alone, while omitting colors and measurements from the data. We use predefined dictionaries of colors (e.g. "red", "white") and measurements (e.g. "small", "XL") to identify appearances in item titles and replace them with constant tokens, one for colors and another for sizes. This baseline simulates an

attempt to tackle the consistency issue by manually identifying few spurious features in the data and hiding them from the model to make it consistent.

4.2.2 CST

We evaluate CST with two configurations, each utilizes a different version of \mathcal{D}_U : 1) the **complete** data (354k groups with 1.3M samples), and 2) **sub-sampled (SS)** data, sampled to be as similar as possible to \mathcal{D}_L 's histogram (yielding 70k groups with 250k samples). Fig. 3 provides an illustrations of those histograms. In order to assign each group of items with one single label, as described earlier, we choose the category with the highest confidence score within the group provided by $f^{\text{base } 1}$.

4.2.3 CGA

In order to empirically evaluate CGA, we construct $\mathcal{D}_{\text{pairs}}$ from \mathcal{D}_U as described earlier and use a T5 model (Raffel et al., 2020) (a large Transformer based seq-2-seq model) as \mathcal{M} , which we fine-tune on $\mathcal{D}_{\text{pairs}}$ for three epochs.

The impact of the filtering score function. We examine two alternatives of the score function s ; 1) BLEU score (Papineni et al., 2002) and 2) a cosine-similarity score that was computed on the output vectors of an all-MinLM-L6-V2 model (All-MinLM-L6-V2). This model maps sentences to

¹Preliminary experiments showed that this method outperformed majority voting.

Original Product Title	Generated Product Title	BLEU score
Polo Ralph Lauren Big Boys Fleece Hoodie	Polo Ralph Lauren Little Boys Fleece Hoodie	0.795
Puff Sleeve T Shirt Ivory Frost	T Shirt	0.135
Blackberries Prepacked 6 Oz	Cranberry Prepacked 6 Oz	0.724
Sunnies Face Airblush in Peached	Sunnies Face Airblush in Peached Wall Poster	0.482
	With Pushpins	
Artistry Signature Color Long-wearing Eye Pencil Brown	Artistry Signature Color Long-wearing Eye Pencil Black	0.850

Table 1: Examples of pairs of original product titles and their corresponding generated ones, together with the computed BLEU score of the pairs.

a 384 dimensional dense vector space and can be used for tasks such as clustering or semantic search. We compute both scores for each pair of original product title and a corresponding generated title. Preliminary experiments show that filtering by the BLEU score results in a more consistent model. For the rest of the experiments we use the BLEU score as s . Table 1 contains some examples of generated titles and their corresponding BLEU score.

Using the T5 model, we generate 8 samples based on each sample from \mathcal{D}_L , and compute the s score of each of those samples. We then perform three experiments to evaluate the impact of the filtering threshold T and the augmentation size N . Results are presented in Figure 4.

The impact of the filtering threshold. For each threshold value $T \in \{0.5, 0.6, 0.7, 0.8\}$, we filter the generated samples. Then, we sub-sample a fixed amount of $N = 200k$ samples into \mathcal{D}_{aug} and train a model on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. As T gets higher, the consistency rate of the trained model increases as well, which indicates the need of a filtering phase.

The impact of the augmentation size. We filter the generated samples using a fixed $T = 0.7$. Out of the remaining generated samples, we sub-sample $N \in \{50k, 100k, 200k, 400k\}$ samples into \mathcal{D}_{aug} , and train a model on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. As N gets higher, the consistency rate of the trained model increases as well, which indicates that adding more generated samples leads to a more consistent model.

The trade off between filtering threshold and augmentation size. We filter the generated samples using different thresholds, and add the filtered samples to \mathcal{D}_{aug} without sub-sampling them. We train a model on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$. Evidently, the consistency rate of the trained model increases when T gets higher but decreases for $T = 0.8$. As T gets higher, the filtered samples are of better quality but

there are fewer of them, reaching an optimal trade off at $T = 0.7$. Thus, for the rest of the paper, we use $T = 0.7$.

5 Results and Discussion

We train each examined model 5 times and present the mean score of the achieved results. For each model, we compare the weighted average F1 score for the accuracy test and the consistency rate of the consistency test. Table 2 presents our results.

The ColorsSizes-Blind model performs similarly to the baseline for both measurements; the slight changes are within the std range, thus making the differences insignificant compared to the baseline model. This is an evidence that the item versioning is more complex than just changing the size or color and includes title rephrasing concepts that are hard to tackle in a trivial way.

In addition, the results show that both of the CST versions, complete and sub-sampled, achieve significantly higher consistency rates than the baseline, gaining lifts of 10% and 7% respectively. On the other hand, both of the methods yield lower F1 scores, reducing lift by 1.65% and 0.6% respectively. A possible cause of the degradation in the F1 score is the differences between the data distribution of \mathcal{D}_L , which we sample the accuracy test from, and the data distribution of \mathcal{D}_U which we use to augment our training data. The fact that using the sub-sampled version of \mathcal{D}_U mitigates most of this degradation supports this claim. An additional cause could be the usage of the noisy pseudo-labels in the augmented set instead of the unavailable ground truth labels. Note that the amount of added data using \mathcal{D}_U to tackle consistency is bigger than the original \mathcal{D}_L , which aims to tackle accuracy. The focus in terms of the training shifts from an accuracy problem to a consistency problem, thus

Method	F1	F1 lift	Cnst. rate	Cnst. lift
Baseline	0.665	-	0.738	-
CS-Blind	0.664	-0.13%	0.740	0.26%
CST-Full	0.654	-1.65%	0.813	10.12%
CST-SS	0.661	-0.6%	0.790	6.99%
CGA	0.667	0.28%	0.771	4.46%

Table 2: Categorization results, indicating the mean. Lift values are all compared to the Baseline model. The std ranges between 0.001 to 0.002 for F1 and 0.001 to 0.009 for the consistency rate.

hurting the F1 of the new model. The higher consistency rate of CST-Full compared to the CST-Sub-Sampled can be explained by a difference of more than 1M samples in the size of \mathcal{D}_{aug} .

Similarly, the CGA method also improves the consistency rate, gaining lift of 4.5%, and doesn't significantly affect the accuracy score. As mentioned, we use a threshold $T = 0.7$, thus including 440k samples in \mathcal{D}_{aug} . These additional samples correspond to a similar distribution as \mathcal{D}_L . The improvement in both the consistency and the accuracy indicates that the generative model is able to correctly learn the real-world item versioning and produce a significant amount of data with high accuracy labels and the same distribution as in the accuracy test.

Summarizing the above, our experiments highlight three key factors in the consistent-SSL framework: 1) **Scale** - enriching the learning set with more examples of item versioning increases the consistency. 2) **Quality** - augmenting the data with real-world samples is better than using generative ones in term of performance. 3) **Distribution** - preserving the original distribution in the augmented set is important for maintaining good accuracy.

6 Conclusions

This work presents a new framework for consistent text categorization in the context of e-Commerce. The aim of this work is to improve a product categorization model that serves various services of a major web company. We address the labeling inconsistency issues found in the categorization of similar items, leading to poor user experience in related recommendation and search applications. Our framework utilizes an unlabeled clustered dataset in two ways: a self-training approach and a generative-augmentation method. We performed a thorough

investigation of the two approaches and investigated several factors that majorly influence their performance. Our experimental results suggest that both proposed methods improve the consistency rate by 4% to 10%, while maintaining the accuracy of the current production model. Finally, our study illustrates the trade off between the quality and the scale of the augmented dataset, and its impact on the performance of both methods.

Limitations

Our work has several limitations. First, our consistency study focuses on our used categorization model and was conducted on only one specific dataset. It might not perfectly generalize to other problems. Second, the proposed solutions are based solely on data augmentation without changing the current production settings and model. Other approaches such as changing the model's objective function to take consistency into account might also benefit the solution. Lastly, in terms of user perspective, while our solution show significant improvement over the baseline, inconsistencies are still visible.

Ethics Statement

This NLP research study was designed and carried out with strict adherence to ethical principles and guidelines. The study was reviewed and approved by our company's research lead prior to the submission. The study followed the ACL conference's guidelines on the use of language data. The researchers take full responsibility for ensuring the ethical conduct of this study and are committed to upholding the highest standards of ethical research practices in NLP.

References

- All-MinLM-L6-V2. 2022. All-minlm-l6-v2. <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>. [Online; accessed 10-October-2022].
- Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.

- M Saiful Bari, Tasnim Mohiuddin, and Shafiq Joty. 2020. Uxla: A robust unsupervised data augmentation framework for zero-resource cross-lingual nlp. *arXiv preprint arXiv:2004.13240*.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and improving consistency in pretrained language models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. *Proceedings of the AAAI conference on artificial intelligence*, 34(05):8018–8025.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2019. Learning the difference that makes a difference with counterfactually-augmented data. *arXiv preprint arXiv:1909.12434*.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. Data augmentation using pre-trained transformer models. *arXiv preprint arXiv:2003.02245*.
- Dong-Hyun Lee et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. *Workshop on challenges in representation learning, ICML*, 3(2):896.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Georgios Rizos, Konstantin Hemker, and Björn Schuller. 2019. Augment to prevent: short-text data augmentation in deep learning for hate-speech classification. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 991–1000.
- Connor Shorten, Taghi M Khoshgoftaar, and Borko Furht. 2021. Text data augmentation for deep learning. *Journal of big Data*, 8(1):1–34.
- Isaac Triguero, Salvador Garcia, and Francisco Herrera. 2015. Self-labeled techniques for semi-supervised learning: taxonomy, software and empirical study. *Knowledge and Information systems*, 42(2):245–284.
- Victor Veitch, Alexander D’Amour, Steve Yadlowsky, and Jacob Eisenstein. 2021. Counterfactual invariance to spurious correlations: Why and how to pass stress tests. *arXiv preprint arXiv:2106.00545*.
- Tianlu Wang, Xuezhi Wang, Yao Qin, Ben Packer, Kang Li, Jilin Chen, Alex Beutel, and Ed Chi. 2020. Catgen: Improving robustness in nlp models via controlled adversarial text generation. *arXiv preprint arXiv:2010.02338*.
- Tianlu Wang, Diyi Yang, and Xuezhi Wang. 2021. Identifying and mitigating spurious correlations for improving robustness in nlp models. *arXiv preprint arXiv:2110.07736*.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. 2020. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268.

A CST Algorithm

Algorithm 1 CST

Input: labeled training data $\mathcal{D}_L = \{(x_i, y_i)\}_{i=1}^l$,
unlabeled data $\mathcal{D}_U = \{X_i\}_{i=l+1}^{l+u}$, $X_i =$
 $\{\hat{x}_j^{(i)}\}_{j=1}^{k_i}$, set function h

- 1: train a base classifier f^{base} on \mathcal{D}_L
- 2: $\mathcal{D}_{\text{aug}} \leftarrow \emptyset$
- 3: **for** $i = l + 1, l + 2, \dots, l + u$ **do**
- 4: $\tilde{y}_i \leftarrow h(X_i; f^{\text{base}})$
- 5: $\mathcal{D}_{\text{aug}} \leftarrow \mathcal{D}_{\text{aug}} \cup \{(\hat{x}_j^{(i)}, \tilde{y}_i)\}_{j=1}^{k_i}$
- 6: train f on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$
- 7: **return** f

otherwise it returns the prediction of f_{i-1} . If the process gets to the end, i.e. f_4 agrees with f_3 on the label until L_3 , it returns the prediction of f_4 as the final prediction.

B CGA Algorithm

Algorithm 2 CGA

Input: labeled training data $\mathcal{D}_L = \{(x_i, y_i)\}_{i=1}^l$,
unlabeled data $\mathcal{D}_U = \{\hat{X}_i\}_{i=l+1}^{l+u}$, $X_i =$
 $\{\hat{x}_j^{(i)}\}_{j=1}^{k_i}$, number of samples to generate
from each original sample n , score function
 s , threshold T

- 1: $\mathcal{D}_{\text{pairs}} = \left\{ \left(\hat{x}_j^{(i)}, \hat{x}_{j'}^{(i)} \right) \mid l + 1 \leq i \leq l + u \wedge j, j' \in [k_i] \right\}$
- 2: train a generative model \mathcal{M} on $\mathcal{D}_{\text{pairs}}$
- 3: $\mathcal{D}_{\text{aug}} \leftarrow \emptyset$
- 4: **for** $i = 1, 2, \dots, l$ **do**
- 5: generate n new samples $\hat{x}_1^{(i)}, \dots, \hat{x}_n^{(i)}$ with \mathcal{M} and x_i
- 6: **for** $j = 1, 2, \dots, n$ **do**
- 7: **if** $s(\hat{x}_j^{(i)}, x_i) \geq T$
- 8: **then** $\mathcal{D}_{\text{aug}} \leftarrow \mathcal{D}_{\text{aug}} \cup \{(\hat{x}_j^{(i)}, y_i)\}$
- 9: train f on $\mathcal{D}_L \cup \mathcal{D}_{\text{aug}}$
- 10: **return** f

C Hierarchical-FastText

Hierarchical-FastText (HFT) consist of 4 FastText models $\{f_i\}_{i=1}^4$. In training time, each f_i is trained over the same data samples, but with different granularity of the labels: f_1 is trained using only the first level of the labels L_1 , f_2 is trained using the first and second levels of the labels L_1 and L_2 and so on. In inference time, we use an iterative method, were at each iteration i for $i = 1, \dots, 4$ we predict the label using f_i . If f_i agrees with f_{i-1} on the label until the level L_{i-1} , the process continues,