

# Atlas: A Dataset and Benchmark for E-commerce Clothing Product Categorization

Venkatesh Umaashankar  
venkatesh.u@ericsson.com  
Ericsson Research  
Chennai, India

Girish Shanmugam S  
s.girishshanmugam@gmail.com  
Uppsala University  
Uppsala, Sweden

Aditi Prakash  
aditi.prakash@colorado.edu  
University of Colorado, Boulder  
Boulder, USA

## ABSTRACT

In e-commerce, it is a common practice to organize the product catalog using product taxonomy. This enables the buyer to easily locate the item they are looking for and also to explore various items available under a category. Product taxonomy is a tree structure with 3 or more levels of depth and several leaf nodes. Product categorization is a large scale classification task that assigns a category path to a particular product. Research in this area is restricted by the unavailability of good real-world datasets and the variations in taxonomy due to the absence of a standard across the different e-commerce stores. In this paper, we introduce a high-quality product taxonomy dataset focusing on clothing products that contain 186,150 images under clothing category with 3 levels and 52 leaf nodes in the taxonomy. We explain the methodology used to collect and label this dataset. Further, we establish the benchmark by comparing image classification and Attention-based Sequence models for predicting the category path. Our benchmark model reaches a micro f-score of 0.92 on the test set. The dataset, code, and pre-trained models are publicly available at <https://github.com/vumaasha/atlas>. We invite the community to improve upon these baselines.

## KEYWORDS

Product Categorization, Attention, Seq to Seq models, Image Classification, Computer Vision

### ACM Reference Format:

Venkatesh Umaashankar, Girish Shanmugam S, and Aditi Prakash. 2020. Atlas: A Dataset and Benchmark for E-commerce Clothing Product Categorization. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom'20)*. ACM, New York, NY, USA, 6 pages.

## 1 INTRODUCTION

With the Internet revolution, e-commerce has become a major platform for selling products to customers. E-commerce stores host a collection of products ranging from electronics to fashion apparel to grocery. A well-organized e-commerce store lets customers navigate through the website with ease and locate the product they are looking for. Unlike a traditional retail store where you can walk in and seek assistance, online retailers rely on their product catalog or categorization to assist shoppers to find their desired product.

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR eCom'20, July 30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s).

Product taxonomy is a tree structure with multiple top and intermediate levels, ending in leaf nodes. Taxonomy classification is the process of assigning a category path to a particular product in the taxonomy tree. E-commerce sites use hierarchical taxonomies to organize products from generic to specific classes where each level provides more specific details about the product than the previous level. For example, *Clothing & Accessories >Men >Winterwear >Sweatshirts & Hoodies*. These classification levels are important for an e-commerce store to perform operations such as search, catalog building, recommendation, classifying user queries [19], which thereby hugely influence customer satisfaction and revenue of e-commerce sites. Currently, most of these product classification mechanisms rely on sellers to provide correct details. Each e-commerce store has its own product taxonomy and a seller typically sells in multiple stores. This implies that the seller has to perform such categorizations manually multiple times. Automating this has the potential benefits of reduced costs and better catalog quality. Our key contributions in this paper are:

- Developed a clean and rich clothing product taxonomy dataset containing 186,150 images and their corresponding product titles which maps to 52 category paths.
- Proposed a methodology to collect large scale product taxonomy dataset which can be easily extended to categories other than Clothing.
- Trained and compared two benchmark models (Image classification and Attention-based Seq to Seq model) that predicts the category path from the product image. Our best model reached an f-score of 0.92 on the test set.
- The dataset, source code, and pre-trained models are made publicly available<sup>1</sup> to encourage future research in this area.

The rest of the paper is organized as follows. Related literature is reviewed in Section 2. In Section 3 we explain the methodology that we used to develop the *Atlas* dataset. In Section 4 we build our benchmark models and explain our model architecture. In Section 5 we provide our training setup and meta parameters to facilitate reproducible research. In Section 6 we summarize our results. Finally, in Section 7 we conclude and provide details about possible directions for future work.

## 2 RELATED WORK

A clean and detailed product taxonomy offers several benefits to both the e-commerce store and its customers. However, creating, maintaining, or adapting an existing categorization standard is not an easy task. Still, most of e-commerce stores want to have

---

<sup>1</sup><https://github.com/vumaasha/Atlas>

flexibility in the way they organize their catalog and create their product taxonomy.

Initially, techniques from information retrieval and machine learning were applied to solve the problem of product categorization. GoldenBullet [4] is a software environment targeted to automatically classify the products, based on their original descriptions and existent classification standards (such as UNSPSC). It integrates different classification algorithms like Vector space model (VSM), K-nearest neighbor and Naive-Bayes classifier algorithms, and some natural language processing techniques to pre-process data. [5] approached product categorization as a hierarchical text classification task. They proposed two different approaches of building separate classifiers for each level in the hierarchy and a flat classifier that directly predicts the leaf level assignment of a document. They used Support Vector Machine (SVM) classifiers for evaluating both the approaches. [11] presented a simple linear classifier based approach for product categorization using mutual information and LDA based features. In general, the computational complexity involved in some of these traditional machine learning techniques is well beyond linear with respect to the number of training examples, features, or classes. The scale of the e-commerce product categorization requires algorithms capable of processing a huge volume of training data in a reasonable time, capable of handling a large number of classes and also capable of making fast real-time predictions [18].

The remarkable progress made in the field of deep learning in recent years has provided a better way to approach this problem. [3] has done a detailed study of using Convolutional Neural Networks (CNN) for the product categorization task. They used the Amazon product dataset provided by [17] and text features such as product titles, navigational breadcrumbs, and list price. [6] used multiple Deep Recurrent Neural Network (RNNs) and generated features from the text metadata. In recent times, Sequential model-based approaches have been widely used for product categorization. [8] modeled product categorization as a Sequence to Sequence Learning, they used product titles which are a sequence of words as input and predicted the category path as a sequence of category levels in the product taxonomy.

Due to the availability of large high-quality image datasets, the field of Image classification [12] has matured a lot in recent times. Noise and ambiguity is a common problem in textual product titles and description. However, most of the e-commerce products tend to have decent product images, this leads to a natural choice of using images for product categorization. [1], [2], [10] and [16] applied computer vision techniques for fashion apparel categorization based on the product images. The closest to our work is by [13] where they use Seq to Seq model with product titles as input to predict category paths using an LSTM Decoder and beam search for inference. We extend their work in this paper by using product images instead of product titles as input for product categorization. Similar to [13], we learn an Attention-based Seq to Seq model.

### 3 ATLAS DATASET

Rakuten made a product classification dataset publicly available in Rakuten Data Challenge [15], However, this dataset contains only the product titles, and the levels in the taxonomy are represented

using numerical IDs instead of plain text. Real-world product taxonomy datasets are not publicly available. Also, there is no widely adopted industry standard for defining product taxonomies. In addition to these, factors like data size, category skewness, and noisy metadata are limiting further research and practical implementation of large scale product categorization. This motivated us to develop a real-world dataset for product categorization.

We developed a new product categorization dataset called *Atlas*. An e-commerce store typically sells products under several top-level categories such as Electronics, Home & Kitchen, Clothing, etc. In this paper, we focused only on clothing products. Our dataset contains data corresponding to 52 products and their title, price, image, and category path.

#### 3.1 Taxonomy Generation

In the e-commerce world, each store has its own taxonomy. For example the category path for 'Jackets' in Flipkart<sup>2</sup> is *Clothing >Men's Clothing >Winter & Seasonal Wear >Jackets* and in Amazon<sup>3</sup> it is *Clothing & Accessories >Men >Jackets*. Treating them as different category paths will lead to noisy taxonomy and training data. We designed our taxonomy based on the similarities in the taxonomy structures across the different e-commerce retailer websites. Our taxonomy is organized to a maximum depth of 3 levels which can assist the consumer to reach their product in not more than 3 clicks. The process of building our taxonomy involved three steps. First, we analyzed and listed the taxonomy structures of popular products, niche, and premium clothing products across different e-commerce stores. Next, we identified the common category paths up to the third level across these websites. Finally, clothing that had the same category paths until the third level were clubbed together irrespective of the dissimilarity in the deeper levels. For example, *Women >Ethnic Wear >Salwar Kameez >Bollywood* and *Women >Ethnic Wear >Salwar Kameez >Anarkali* are grouped under the category *Women >Ethnic Wear >Salwar Kameez* as they have similar category paths until the third level beyond which it branches into different nodes. Our final taxonomy tree is not an exhaustive list of all clothing categories but covers popular Western Wear and niche Ethnic Wear especially from the Indian Clothing collection. Each of the categories in the taxonomy tree has a maximum depth of 3 levels and totals to 52 category paths. Our taxonomy tree and a few sample products from some of the categories in our dataset can be found here <sup>4</sup>.


#### 3.2 Data Collection

We crawled the product listings from popular Indian e-commerce stores. We manually created a mapping of the store's category path that maps to a category path in our taxonomy. We used web scraping tools Scrapy and Selenium. The crawlers extract the information from the HTML content of the product page using CSS selectors. We extracted the *product title, breadcrumb, image, and price* corresponding to each product in the product listings. The attributes extracted are stored in JSON format.

<sup>2</sup><https://www.flipkart.com/>

<sup>3</sup><https://www.amazon.in/>

<sup>4</sup><https://github.com/vumaasha/Atlas/tree/master/dataset#11-taxonomy-generation>



**Figure 1: Examples of (a) Zoomed (dirty) and (b) Normal (clean) images from our *Atlas* dataset. The Zoomed images show close-ups of the apparel or cropped versions of the image that make it difficult to recognize the product, whereas the Normal images show figures with the entire product visible.**

### 3.3 Data Cleaning

It is typical for an e-commerce store to show several images for a single product. Out of these images, not all the images are necessarily a good representative image of the product. Some images might display packaging, installation instructions, etc. In the case of clothing, we found that many product listings also included zoomed-in images that display intrinsic details such as the texture of the fabric, brand labels, button, and pocket styles as shown in Figure 1. Without the context of the product listing, it would be even hard for a human to identify the corresponding product. Including these zoomed-in images would drastically affect the quality of the dataset. To find and remove these noisy images manually would take considerable time and effort. We modeled this as a binary classification task (Zoomed Vs Normal Images) and compared Linear SVM with a simple 3 layer CNN (Figure 2) based classification models. We prepared the training data by visual inspection. We segregated noisy and high-quality images into two different folders by looking at the thumbnails of hundreds of product images in a go. Our models were trained on 6005 normal images and 1054 zoomed images and the performance metrics on the test are shown in the Table. 1. We used computer vision-based features such as contours and histogram of gradients as input for our LinearSVM. We automated the process of filtering out the noisy images using the CNN model due to its superior performance compared to that of LinearSVM.

**Table 1: Metrics for the models used to predict Zoomed Vs Normal images**

	CNN			SVM		
	precision	recall	f-score	precision	recall	f-score
Normal	0.99	0.99	0.99	0.91	0.99	0.95
Zoomed	0.95	0.95	0.95	0.86	0.48	0.62
Average	0.98	0.98	0.98	0.91	0.91	0.90

## 4 BENCHMARK MODELS FOR PRODUCT CATEGORIZATION

### 4.1 Resnet34 based Image Classification

We use the `cnn_learner` available in `fast.ai` implementation to train our Image classification model. This uses Resnet34 architecture

as the backbone of the model, which is followed by AdaptiveConcatPool2d, Flatten, and 2 blocks of `[nn.BatchNorm1d, nn.Dropout, nn.Linear, nn.ReLU]` layers. The first block will have several inputs inferred from the backbone arch Resnet34 (512) and the second one will have several outputs equal to the number of classes (52) without `nn.ReLU` activation.

### 4.2 Attention based Seq to Seq Model

We approach the product categorization problem as a sequence prediction problem by leveraging the dependency between each level in the category path. We use Attention-based Encoder-Decoder Neural Network architecture to generate sequences. The Encoder is a 101 layered Residual Network(ResNet) trained on the ImageNet classification task which converts the input image to a fixed size vector. The Decoder is a combination of Long Short-Term Memory(LSTM) along with Attention Network which combines the Encoder output and Attention weights to predict category paths as sequences. Our architecture is similar to works by [22] and [23] used for Neural Machine Translation and image captioning respectively. We extended the source code from the pytorch tutorial to image captioning repository by Sagar Vinodababu<sup>5</sup>. Figure 3 shows some of the category paths generated by our model on test images that are not seen during training or validation. From this figure, it can be clearly seen that to generate each category level our model focuses on different parts of the image. To predict the first category level, which is the gender, our model has focused on the face and it has focused on the actual region of the clothing products to predict the next category levels, 'SareeBlouse' and 'Kurta'.

**4.2.1 Encoder and Decoder.** In Encoder, we use Convolutional Neural Network (CNN) to produce fixed size vectors. The images in *Atlas* dataset have different dimensions as they were collected from different sources. All these images are resized to have a uniform dimension of 150\*150 pixels before being fed as input to the Encoder. The input images are then represented by the 3 color channels of RGB values. The Encoder uses a 101 layered Residual Network pre-trained on the ImageNet classification task which is shown in Figure 4. As we use the Encoder only to encode images and not for classifying them, we remove the last two layers (linear and pooling layers) from the ResNet-101 model proposed by [7]. The images are resized to fixed size by adding a 2D adaptive average pooling layer which enables the Encoder to accept images of variable sizes. The final encoding produced by the Encoder will have the dimensions: `batch_size, 14, 14, 2048`.

Recurrent Neural Networks(RNN) are popular for sequential classification task as it considers both the current input and the learnings from the previously received inputs for prediction. Usually, RNN's have short term memory but when combined with Long Short-Term Memory (LSTM) Network they have long term memory as LSTM's contain their information in a memory. We have a stacked LSTM Network along with Attention in our Decoder which is shown in Figure 4.

The Attention Network shown in Figure 4, learns which part of the image has to be focused to predict the next level in the category path while performing the sequence classification task. The

<sup>5</sup><https://github.com/sgrvinod/a-PyTorch-Tutorial-to-Image-Captioning>




Figure 2: Architecture of Zoomed Vs Normal Model

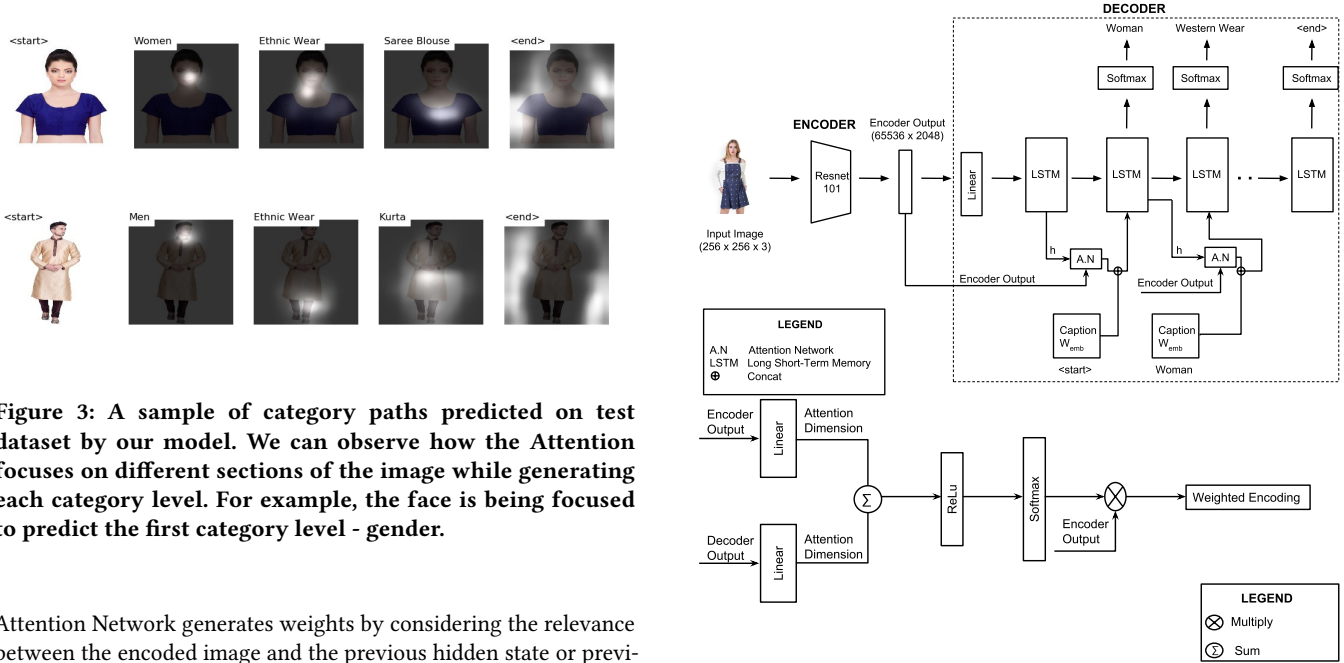


Figure 3: A sample of category paths predicted on test dataset by our model. We can observe how the Attention focuses on different sections of the image while generating each category level. For example, the face is being focused to predict the first category level - gender.

Attention Network generates weights by considering the relevance between the encoded image and the previous hidden state or previous output of the Decoder. It consists of linear layers that transform the encoded image and the previous Decoder’s output to the same size. These vectors are summed together and passed to another linear layer. This layer calculates the values to be Softmaxed and then passes the values to a ReLU layer. A final softmax layer calculates the weights  $\alpha$ s of the pixels which add up to 1. If there are  $P$  pixels in our encoded image, then at each time step  $t$ ,

$$\sum_p^P \alpha_{p,t} = 1 \quad (1)$$

We use a weighted average across all the pixels instead of a simple average so that the important pixels are assigned greater weights.

The Decoder receives the encoded image from the Encoder using which it initializes the hidden and cell state of the LSTM model through two linear layers. Two virtual category levels  $\langle \text{start} \rangle$  and  $\langle \text{end} \rangle$  which denote the beginning and end of the sequence are added to the category path. The Decoder LSTM uses teacher forcing proposed by [21] for training. The Decoder uses a  $\langle \text{start} \rangle$  marker which is considered to be the zeroth category level. The  $\langle \text{start} \rangle$  marker along with the encoded image is used to generate the first-top level of the category path. Subsequently, all other levels are

Figure 4: Encoder - Decoder with Attention Network

predicted using the sequence generated so far along with the Attention weights. An  $\langle \text{end} \rangle$  marker is used to mark the end of a category path. The Decoder stops decoding the sequence further as soon it generates the  $\langle \text{end} \rangle$  marker. At each time step, the Decoder computes the weights and Attention weighted encoding from the Attention Network using its previous hidden state. Another linear layer is added to create a sigmoid-activated gate and the Attention weighted encodings are passed through it and concatenated with the embedding of the previously generated category path and fed into the LSTM Decoder to generate the new hidden state which is also the next predicted level. The next level is predicted using a final softmax layer from the hidden state of the Decoder. The softmax layer transforms the hidden state into scores which are stored for further utilization in beam search for selecting 'k' best levels.

## 5 TRAINING

### 5.1 Model hyperparameters

**5.1.1 Zoomed Vs Normal LinearSVM Model.** We trained LinearSVM available in Scikit-learn with C set to 0.0001, class weight set to 'balanced' using hinge loss. The optimal C value was identified using a grid search.

**5.1.2 Zoomed Vs Normal CNN Model.** We trained for 10 epochs using Binary CrossEntropy as loss function, RMSProp Optimizer with a learning rate set to 0.001, rho set to 0.9, and decay set to 0.0.

**5.1.3 Resnet34 based Image Classification.** We trained for 17 epochs using Categorical CrossEntropy as loss function and Leslie Smith's one cycle policy [20] for choosing the learning rate. We used early stopping to terminate the training process when the decrease in validation loss is less than 0.001 for 3 consecutive epochs.

**5.1.4 Attention based Seq to Seq Model.** We trained our model in GPU for 3 epochs with a batch size of 128 and a dropout rate as 0.5 after which the validation accuracy stop improving. We used Adam optimizers with a learning rate of 1e-4 and 4e-4 for Encoder and Decoder respectively. We picked the beam width as 5 based on our experiments. The regularization parameter for doubly stochastic attention was set to 1 and gradient clipping was set to an absolute value of 5. The pre-trained model can be downloaded from here<sup>6</sup>.

### 5.2 Hardware


- Nvidia GPU GEFORCE GTX 1080 Ti 11GB RAM
- Intel® Xeon® Processor E5-2650 v4 30M Cache, 2.20 GHz, 12 Cores, 24 Threads
- 250 GB RAM
- CentOS 7

## 6 RESULTS

We evaluated the proposed model on our dataset having 186,150 clothing images and their category paths. We split our dataset into train, validation, and test sets similar to the splits used in the work by [9]. Stratified random sampling was carried out on our dataset with training set having 65% of data(119,155 images), 5% in the validation set(11,147 images), and 30% in the test set(55,848 images). The Resnet34 classification model and the Seq to Seq model trained on our Atlas dataset achieved an overall micro f-score of 92% and 90% respectively. A comparison of the f-scores of both the benchmark models over support size of leaf categories is shown in Figure 5. Though we observe that the classification model's performance is better than Seq to Seq model, we believe the reason is that we have only 52 categories at the moment. As the number of categories increases, the structure in the taxonomy can be leveraged better using Seq to Seq model. In addition to Seq to Seq models predicting the category paths, it also explains the reason behind the predictions which is shown in Figure 3.

[14] claims that using Seq to Seq model for product categorization helps to identify new category paths in the taxonomy. However, in our experiments, we have observed that all the new category paths that are generated by the Seq to Seq model are not always valid. In our case our model generated 5 new category paths which

<sup>6</sup><https://goo.gl/forms/C1824kmbuVo7H6H3>



**Figure 5: F-scores of our benchmark models over leaf level categories ordered by their sample size. Note that the sample size in x axis is in log scale**

are shown in Table 2 out of which we found only 2 to be valid. Therefore, a manual inspection of newly created category paths is needed to filter out the category paths which could be used to enrich the taxonomy.

**Table 2: Valid and invalid category paths created by Seq to Seq model**

Valid Category Paths	Invalid Category Paths
Women >Western Wear >Jackets	Men >Western Wear >Dresses
Women >Western Wear >Blazers&Suits	Men >Western Wear >Tanktops&Camisoles
	Women >Inner Wear >Shorts

## 7 CONCLUSION AND FUTURE WORKS

This paper introduces *Atlas*, a fashion apparel dataset with 186,150 apparel images along with their corresponding product titles. We have open-sourced the codebase and the procedure to build the dataset of images and their taxonomy. We have proposed two benchmark models using classification and Attention-based Sequence approaches to predict product taxonomy.

In the future, we plan to extend our *Atlas* dataset by adding more categories and products thereby increasing the total number of category paths. We would avoid the generation of invalid category paths in our Seq to Seq model by considering the taxonomy structure while decoding and explore Transformer Networks instead of Recurrent Neural Networks (RNN).

## 8 ACKNOWLEDGEMENTS

The First Author Venkatesh Umaashankar worked extensively in the problem of Product Categorization using text attributes during his tenure at Indix. He thanks Krishna Sangeeth, Sriram Ramachandrasekaran, Anirudh Venkataraman, Manoj Mahalingam, Rajesh Muppalla, and Sridhar Venkatesh for their help and support.

## REFERENCES

- [1] Lukas Bossard, Matthias Dantone, Christian Leistner, Christian Wengert, Till Quack, and Luc Van Gool. 2012. Apparel classification with style. In *Asian conference on computer vision*. Springer, 321–335.
- [2] Huizhong Chen, Andrew Gallagher, and Bernd Girod. 2012. Describing clothing by semantic attributes. In *European conference on computer vision*. Springer, 609–623.

- [3] Pradipto Das, Yandi Xia, Aaron Levine, Giuseppe Di Fabbrizio, and Ankur Datta. 2017. Web-scale language-independent cataloging of noisy product listings for e-commerce. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, Vol. 1. 969–979.
- [4] Ying Ding, M Korotkiy, Borys Omelayenko, V Kartseva, V Zykov, Michel Klein, Ellen Schulten, and Dieter Fensel. 2002. Goldenbullet: Automated classification of product data in e-commerce. In *Proceedings of the 5th international conference on business information systems*.
- [5] Susan Dumais and Hao Chen. 2000. Hierarchical classification of Web content. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 256–263.
- [6] Jung-Woo Ha, Hyuna Pyo, and Jeonghee Kim. 2016. Large-scale item categorization in e-commerce using multiple recurrent neural networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 107–115.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [8] Makoto Hiramatsu and Kei Wakabayashi. 2018. Encoder-Decoder Neural Networks for Taxonomy Classification. In *eCOM@SIGIR (CEUR Workshop Proceedings)*, Vol. 2319. CEUR-WS.org.
- [9] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3128–3137.
- [10] M Hadi Kiapour, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. 2014. Hipster wars: Discovering elements of fashion styles. In *European conference on computer vision*. Springer, 472–488.
- [11] Zornitsa Kozareva. 2015. Everyone likes shopping! multi-class product categorization for e-commerce. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1329–1333.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [13] Maggie Yundi Li, Stanley Kok, and Stanley Kok. 2018. Unconstrained Product Categorization with Sequence-to-Sequence Models. In *eCOM@SIGIR (CEUR Workshop Proceedings)*, Vol. 2319. CEUR-WS.org.
- [14] Maggie Yundi Li, Stanley Kok, and Liling Tan. 2018. Don’t Classify, Translate: Multi-Level E-Commerce Product Categorization Via Machine Translation. *arXiv preprint arXiv:1812.05774* (2018).
- [15] Yiu-Chang Lin, Pradipto Das, and Ankur Datta. 2018. Overview of the SIGIR 2018 eCom Rakuten Data Challenge. In *eCOM@SIGIR (CEUR Workshop Proceedings)*, Vol. 2319. CEUR-WS.org.
- [16] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1096–1104.
- [17] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 43–52.
- [18] Dan Shen, Jean-David Ruvini, and Badrul Sarwar. 2012. Large-scale item categorization for e-commerce. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. ACM, 595–604.
- [19] Michael Skinner and Surya Kallumadi. 2019. E-commerce Query Classification Using Product Taxonomy Mapping: A Transfer Learning Approach.. In *eCOM@SIGIR*.
- [20] Leslie N Smith. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 464–472.
- [21] Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation* 1, 2 (1989), 270–280.
- [22] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [23] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*. 2048–2057.