



Multimodal Retrieval in E-Commerce

From Categories to Images, Text, and Back

Mariya Hendriksen^(✉)

Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands
m.hendriksen@uva.nl

Abstract. E-commerce provides rich multimodal data that is barely leveraged in practice. The majority of e-commerce search mechanisms are uni-modal, which are cumbersome and often fail to grasp the customer's needs. For the Ph.D. we conduct research aimed at combining information across multiple modalities to improve search and recommendations in e-commerce. The research plans are organized along the two principal lines. First, motivated by the mismatch between a textual and a visual representation of a given product category, we propose the task of category-to-image retrieval, i.e., the problem of retrieval of an image of a category expressed as a textual query. Besides, we propose a model for the task. The model leverages information from multiple modalities to create product representations. We explore how adding information from multiple modalities impacts the model's performance and compare our approach with state-of-the-art models. Second, we consider fine-grained text-image retrieval in e-commerce. We start off by considering the task in the context of reproducibility. Moreover, we address the problem of attribute granularity in e-commerce. We select two state-of-the-art (SOTA) models with distinct architectures, a CNN-RNN model and a Transformer-based model, and consider their performance on various e-commerce categories as well as on object-centric data from general domain. Next, based on the lessons learned from the reproducibility study, we propose the model for the fine-grained text-image retrieval.

1 Motivation

Multimodal retrieval is an important but understudied problem in e-commerce [48]. Even though e-commerce products are associated with rich multimodal information, research currently focuses mainly on textual and behavioral signals to support product search and recommendation [1, 15, 42]. The majority of prior work in multimodal retrieval for e-commerce focuses on applications in the fashion domain, such as recommendation of fashion items [34] and cross-modal fashion retrieval [13, 25]. In the more general e-commerce domain, multimodal retrieval has not been explored that well yet [17, 31]. Motivated by the knowledge gap, we lay out two directions for the research agenda: category-to-image retrieval, and fine-grained text-image retrieval (Fig. 1).




Fig. 1. Dissertation overview.

Category-to-Image Retrieval. First, we focus on the category information in e-commerce. Product category trees are a key component of modern e-commerce as they assist customers when navigating across large product catalogues [16, 24, 46, 50]. Yet, the ability to retrieve an image for a given product category remains a challenging task mainly due to noisy category and product data, and the size and dynamic character of product catalogues [28, 48]. Motivated by this challenge, we introduce the task of retrieving a ranked list of relevant images of products that belong to a given category, which we call the *category-to-image* (CtI) retrieval task. Unlike image classification tasks that operate on a predefined set of classes, in the CtI retrieval task we want to be able not only to understand which images belong to a given category but also to generalize towards unseen categories. Use cases that motivate the CtI retrieval task include (1) the need to showcase different categories in search and recommendation results [24, 46, 48]; (2) the task can be used to infer product categories in the cases when product categorical data is unavailable, noisy, or incomplete [52]; and (3) the design of cross-categorical promotions and product category landing pages [39].

Fine-Grained Text-Image Retrieval. Second, we address the problem of fine-grained text-image retrieval. Text-image retrieval is the task of finding similar items across textual and visual modalities. Successful performance on the task depends on the domain. In the general domain, where images typically depict complex scenes of objects in their natural contexts information across modalities is matched coarsely. Some examples of such datasets include MS COCO [33], and Flickr30k [53]. By contrast, in the e-commerce domain, where there is typically one object per image, fine-grained matching is more important. Therefore, we focus on *fine-grained text-image retrieval*. We define the task as a combination of two subtasks: 1. *text-to-image retrieval*: given a noun phrase that describes an object, retrieve the image that depicts to the object; 2. *image-to-text retrieval*: given an image of an object, retrieve the noun phrase that describes an object.

We start off by examining the topic in the context of reproducibility. Reproducibility is one of the major pillars of the scientific method and is of utmost

importance for Information Retrieval (IR) as a discipline rooted in experimentation [10]. One of the first works that touch upon reproducibility in IR is the study by Armstrong et al. [2] where the authors conducted a longitudinal analysis of papers published in proceedings of CIKM and SIGIR between 1998–2008 and discovered that the ad-hoc retrieval was not measurably improving. Later on, Yang et al. [51] provided a meta-analysis of results reported on the TREC Robust04 and found out that some of the more recent neural models were outperformed by strong baselines. Similar discoveries were made in the domain of recommender systems research [5,6]. Motivated by the findings, we explore the reproducibility of fine-grained text-image retrieval results. More specifically, we examine how SOTA models for fine-grained text-image fashion retrieval generalize towards other categories of e-commerce products. After analyzing SOTA models in the domain, we plan to improve upon them in a subsequent future work.

2 Related Work

Category-to-Image Retrieval. Early work in image retrieval grouped images into a restricted set of semantic categories and allowed users to retrieve images by using category labels as queries [44]. Later work allowed for a wider variety of queries ranging from natural language [20,49], to attributes [37], to combinations of multiple modalities (e.g., title, description, and tags) [47]. Across these multimodal image retrieval approaches we find three common components: (1) an image encoder, (2) a query encoder, and (3) a similarity function to match the query to images [14,40]. Depending on the focus of the work some components might be pre-trained, whereas the others are optimized for a specific task. In our work, we rely on pre-trained image and text encoders but learn a new multimodal composite of the query to perform CtI retrieval.

Fine-Grained Text-Image Retrieval. Early approaches to cross-modal mapping focused on correlation maximization through canonical correlation analysis [18,19,45]. Later approaches centered around convolutional and recurrent neural networks [11,22,23,29]. They were further expanded by adding attention on top of encoders [29,35,38]. More recently, inspired by the success of transformers [8], a line of work centered around creating a universal vision-language encoder emerged [4,30,32,36]. To address the problem of attribute granularity in the context of cross-modal retrieval, a line of work proposed to segment images into fragments [27], use attention mechanisms [26], combine image features across multiple levels [13], use pre-trained BERT as a backbone [12,54]. Unlike prior work in this domain that focused on fashion, we focus on the general e-commerce domain.

3 Research Description and Methodology

The dissertation comprises two parts. Below, we describe every part of the thesis and elaborate on the methodology.

Category-to-Image Retrieval. Product categories are used in various contexts in e-commerce. However, in practice, during a user’s session, there is often a mismatch between a textual and a visual representation of a given category. Motivated by the problem, we introduce the task of category-to-image retrieval in e-commerce and propose a model for the task.

We use the XMarket dataset recently introduced by Bonab et al. [3] that contains textual, visual, and attribute information of e-commerce products as well as a category tree. Following [7, 21, 43] we use BM25, MPNet, CLIP as our baselines. To evaluate model performance, we use Precision@K where $K = \{1, 5, 10\}$, mAP@K where $K = \{5, 10\}$, and R-precision.

RQ1.1 *How do baseline models perform on the CtI retrieval task? Specifically, how do unimodal and bi-modal baseline models perform? How does the performance differ w.r.t. category granularity?*

To answer the question, we feed BM25 corpora that contain textual product information, i.e., product titles. We use an MPNet in a zero-shot manner. For all the products in the dataset, we pass the product title through the model. During the evaluation, we pass a category expressed as textual query through MPNet and retrieve top- k candidates ranked by cosine similarity w.r.t. the target category. We compare categories of the top- k retrieved candidates with the target category. Besides, we use pre-trained CLIP in a zero-shot manner with a text transformer and a vision transformer (ViT) [9] configuration. We pass the product image through the image encoder. For evaluation, we pass a category through the text encoder and retrieve top- k image candidates ranked by cosine similarity w.r.t. the target category. We compare categories of the top- k retrieved image candidates with the target category.

RQ1.2 *How does a model, named CLIP-I, that uses product image information for building product representations impact the performance on the CtI retrieval task?*

To answer the question, we build product representations by training on e-commerce data. We investigate how using product image data for building product representations impacts performance on the CtI retrieval task. To introduce visual information, we extend CLIP in two ways: (1) We use ViT from CLIP as an image encoder. We add a product projection head that takes as an input product visual information. (2) We use the text encoder from MPNet as category encoder; we add a category projection head on top of the category encoder. We name the resulting model CLIP-I. We train CLIP-I on category-product pairs from the training set. We only use visual information for building product representations.

RQ1.3 *How does CLIP-IA, which extends CLIP-I with product attribute information, perform on the CtI retrieval task?*

To answer the question, we extend CLIP-I by introducing attribute information to the product information encoding pipeline. We add an attribute encoder through which we obtain a representation of product attributes. We concatenate the resulting attribute representation with image representation and pass the resulting vector to the product projection head. Thus, the resulting product

representation \mathbf{p} is based on both visual and attribute product information. We name the resulting model CLIP-IA. We train CLIP-IA on category-product pairs and we use visual and attribute information for building product representation.

RQ1.4 *And finally, how does CLIP-ITA, which extends CLIP-IA with product text information, perform on the CtI task?*

To answer the question, we investigate how extending the product information processing pipeline with the textual modality impacts performance on the CtI retrieval task. We add a title encoder to the product information processing pipeline and use it to obtain title representation. We concatenate the resulting representation with product image and attribute representations. We pass the resulting vector to the product projection head. The resulting model is CLIP-ITA. We train and test CLIP-ITA on category-product pairs. We use visual, attribute, and textual information for building product representations. The results are to be published in ECIR'22 [16]. The follow-up work is planned to be published at SIGIR 2023.

Fine-Grained Text-Image Retrieval. The ongoing work is focused on fine-grained text-image retrieval in the context of reproducibility. For the experiments, we select two SOTA models for fine-grained cross-modal fashion retrieval, each model with distinctive architecture. One of them is based on Transformer while another one is CNN-RNN-based. The Transformer-based model is Kaleido-BERT [54], that extends BERT [8]. Another model is a Multi-level Feature approach (MLF) [13]. Both models claim to deliver SOTA performance by being able to learn image representations that can better represent fine-grained attributes. They were evaluated on Fashion-Gen dataset [41] but, to the best of our knowledge, were not compared against each other.

In the work, we aim to answer the following research questions:

RQ2.1 *How well Kaleido-BERT and MLF perform on data from an e-commerce category that is different from Fashion?*

RQ2.2 *How well both models generalize beyond e-commerce domain? More specifically, how do they perform on object-centric data from the general domain?*

RQ2.3 *How Kaleido-BERT and MLF compare to each other w.r.t performance?*

The results are planned to be published as a paper at SIGIR 2022. The follow-up work is planned to be published at ECIR 2023.

References

1. Ariannezhad, M., Jullien, S., Nauts, P., Fang, M., Schelter, S., de Rijke, M.: Understanding multi-channel customer behavior in retail. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management, pp. 2867–2871 (2021)
2. Armstrong, T.G., Moffat, A., Webber, W., Zobel, J.: Improvements that don't add up: ad-hoc retrieval results since 1998. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, pp. 601–610. Association for Computing Machinery (2009)

3. Bonab, H., Aliannejadi, M., Vardasbi, A., Kanoulas, E., Allan, J.: Cross-market product recommendation. In: CIKM. ACM (2021)
4. Chen, Y.C., et al.: UNITER: learning universal image-text representations. arXiv preprint [arXiv:1909.11740](https://arxiv.org/abs/1909.11740) (2019)
5. Dacrema, M.F., Cremonesi, P., Jannach, D.: Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In: Proceedings of the 13th ACM Conference on Recommender Systems, pp. 101–109 (2019)
6. Dacrema, M.F., Boglio, S., Cremonesi, P., Jannach, D.: A troubling analysis of reproducibility and progress in recommender systems research. *ACM Trans. Inf. Syst. (TOIS)* **39**(2), 1–49 (2021)
7. Dai, Z., Lai, G., Yang, Y., Le, Q.V.: Funnel-transformer: filtering out sequential redundancy for efficient language processing. arXiv preprint [arXiv:2006.03236](https://arxiv.org/abs/2006.03236) (2020)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
9. Dosovitskiy, A., et al.: An image is worth 16x16 words: transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) (2020)
10. Ferro, N., Fuhr, N., Järvelin, K., Kando, N., Lippold, M., Zobel, J.: Increasing reproducibility in IR: findings from the Dagstuhl seminar on “reproducibility of data-oriented experiments in e-science”. In: ACM SIGIR Forum, vol. 50, pp. 68–82. ACM New York (2016)
11. Frome, A., et al.: Devise: a deep visual-semantic embedding model. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems* 26, pp. 2121–2129. Curran Associates Inc (2013)
12. Gao, D., et al.: FashionBERT: text and image matching with adaptive loss for cross-modal retrieval. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2251–2260 (2020)
13. Goei, K., Hendriksen, M., de Rijke, M.: Tackling attribute fine-grainedness in cross-modal fashion search with multi-level features. In: SIGIR 2021 Workshop on eCommerce. ACM (2021)
14. Gupta, T., Vahdat, A., Chechik, G., Yang, X., Kautz, J., Hoiem, D.: Contrastive learning for weakly supervised phrase grounding. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) *ECCV 2020*. LNCS, vol. 12348, pp. 752–768. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58580-8_44
15. Hendriksen, M., Kuiper, E., Nauts, P., Schelter, S., de Rijke, M.: Analyzing and predicting purchase intent in e-commerce: anonymous vs. identified customers. arXiv preprint [arXiv:2012.08777](https://arxiv.org/abs/2012.08777) (2020)
16. Hendriksen, M., Bleeker, M., Vakulenko, S., van Noord, N., Kuiper, E., de Rijke, M.: Extending CLIP for category-to-image retrieval in e-commerce. In: Hagen, M., et al. (eds.) *ECIR 2022*. LNCS, vol. 13186, pp. 289–303. Springer, Cham (2022)
17. Hewawalpita, S., Perera, I.: Multimodal user interaction framework for e-commerce. In: 2019 International Research Conference on Smart Computing and Systems Engineering (SCSE), pp. 9–16. IEEE (2019)
18. Hodosh, M., Young, P., Hockenmaier, J.: Framing image description as a ranking task: data, models and evaluation metrics. *J. Artif. Intell. Res.* **47**, 853–899 (2013)
19. Hotelling, H.: Relations between two sets of variates. In: Kotz, S., Johnson, N.L. (eds.) *Breakthroughs in Statistics*. SSS, pp. 162–190. Springer, New York (1992). https://doi.org/10.1007/978-1-4612-4380-9_14

20. Hu, R., Xu, H., Rohrbach, M., Feng, J., Saenko, K., Darrell, T.: Natural language object retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4555–4564 (2016)
21. Jabeur, L.B., Soulier, L., Tamine, L., Mousset, P.: A product feature-based user-centric ranking model for e-commerce search. In: Fuhr, N., et al. (eds.) CLEF 2016. LNCS, vol. 9822, pp. 174–186. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44564-9_14
22. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3128–3137 (2015)
23. Kiros, R., Salakhutdinov, R., Zemel, R.S.: Unifying visual-semantic embeddings with multimodal neural language models. arXiv preprint [arXiv:1411.2539](https://arxiv.org/abs/1411.2539) (2014)
24. Kondylidis, N., Zou, J., Kanoulas, E.: Category aware explainable conversational recommendation. arXiv preprint [arXiv:2103.08733](https://arxiv.org/abs/2103.08733) (2021)
25. Laenen, K., Moens, M.-F.: Multimodal neural machine translation of fashion e-commerce descriptions. In: Kalbaska, N., Sádaba, T., Cominelli, F., Cantoni, L. (eds.) FACTUM 2019, pp. 46–57. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-15436-3_4
26. Laenen, K., Moens, M.F.: A comparative study of outfit recommendation methods with a focus on attention-based fusion. *Inf. Process. Manag.* **57**(6), 102316 (2020)
27. Laenen, K., Zoghbi, S., Moens, M.F.: Cross-modal search for fashion attributes. In: Proceedings of the KDD 2017 Workshop on Machine Learning Meets Fashion, vol. 2017, pp. 1–10. ACM (2017)
28. Laenen, K., Zoghbi, S., Moens, M.F.: Web search of fashion items with multimodal querying. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 342–350 (2018)
29. Lee, K.H., Chen, X., Hua, G., Hu, H., He, X.: Stacked cross attention for image-text matching. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 201–216 (2018)
30. Li, G., Duan, N., Fang, Y., Jiang, D., Zhou, M.: Unicoder-VL: a universal encoder for vision and language by cross-modal pre-training. arXiv preprint [arXiv:1908.06066](https://arxiv.org/abs/1908.06066) (2019)
31. Li, H., Yuan, P., Xu, S., Wu, Y., He, X., Zhou, B.: Aspect-aware multimodal summarization for Chinese e-commerce products. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 8188–8195 (2020)
32. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: VisualBERT: a simple and performant baseline for vision and language. arXiv preprint [arXiv:1908.03557](https://arxiv.org/abs/1908.03557) (2019)
33. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
34. Lin, Y., Ren, P., Chen, Z., Ren, Z., Ma, J., de Rijke, M.: Improving outfit recommendation with co-supervision of fashion generation. In: The World Wide Web Conference, pp. 1095–1105 (2019)
35. Liu, C., Mao, Z., Liu, A.A., Zhang, T., Wang, B., Zhang, Y.: Focus your attention: a bidirectional focal attention network for image-text matching. In: Proceedings of the 27th ACM International Conference on Multimedia, pp. 3–11 (2019)
36. Lu, J., Batra, D., Parikh, D., Lee, S.: ViBERT: pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: Advances in Neural Information Processing Systems, pp. 13–23 (2019)

37. Nagarajan, T., Grauman, K.: Attributes as operators: factorizing unseen attribute-object compositions. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 169–185 (2018)
38. Nam, H., Ha, J.W., Kim, J.: Dual attention networks for multimodal reasoning and matching. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 299–307 (2017)
39. Nielsen, J., Molich, R., Snyder, C., Farrell, S.: E-commerce user experience. Nielsen Norman Group (2000)
40. Radford, A., et al.: Learning transferable visual models from natural language supervision. arXiv preprint [arXiv:2103.00020](https://arxiv.org/abs/2103.00020) (2021)
41. Rostamzadeh, N., et al.: Fashion-Gen: the generative fashion dataset and challenge. arXiv preprint [arXiv:1806.08317](https://arxiv.org/abs/1806.08317) (2018)
42. Rowley, J.: Product search in e-shopping: a review and research propositions. *J. Consum. Market.* (2000)
43. Shen, S., et al.: How much can clip benefit vision-and-language tasks? arXiv preprint [arXiv:2107.06383](https://arxiv.org/abs/2107.06383) (2021)
44. Smeulders, A., Worring, M., Santini, S., Gupta, A., Jain, R.: Content-based image retrieval at the end of the early years. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(12), 1349–1380 (2000)
45. Socher, R., Fei-Fei, L.: Connecting modalities: semi-supervised segmentation and annotation of images using unaligned text corpora. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 966–973. IEEE (2010)
46. Tagliabue, J., Yu, B., Beaulieu, M.: How to grow a (product) tree: personalized category suggestions for ecommerce type-ahead. arXiv preprint [arXiv:2005.12781](https://arxiv.org/abs/2005.12781) (2020)
47. Thomee, B., et al.: YFCC100M: the new data in multimedia research. *Commun. ACM* **59**(2), 64–73 (2016)
48. Tsagakias, M., King, T.H., Kallumadi, S., Murdock, V., de Rijke, M.: Challenges and research opportunities in ecommerce search and recommendations. In: SIGIR Forum, vol. 54, no. 1 (2020)
49. Vo, N., et al.: Composing text and image for image retrieval—an empirical odyssey. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6439–6448 (2019)
50. Wirojwatanakul, P., Wangperawong, A.: Multi-label product categorization using multi-modal fusion models. arXiv preprint [arXiv:1907.00420](https://arxiv.org/abs/1907.00420) (2019)
51. Yang, W., Lu, K., Yang, P., Lin, J.: Critically examining the “neural hype”: weak baselines and the additivity of effectiveness gains from neural ranking models. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1129–1132. Association for Computing Machinery (2019)
52. Yashima, T., Okazaki, N., Inui, K., Yamaguchi, K., Okatani, T.: Learning to describe e-commerce images from noisy online data. In: Lai, S.-H., Lepetit, V., Nishino, K., Sato, Y. (eds.) ACCV 2016. LNCS, vol. 10115, pp. 85–100. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-54193-8_6
53. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: new similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguist.* **2**, 67–78 (2014)
54. Zhuge, M., et al.: Kaleido-BERT: vision-language pre-training on fashion domain. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12647–12657 (2021)