

Tilburg University

## SIGIR 2021 E-Commerce Workshop Data Challenge

Tagliabue, Jacopo; Greco, Ciro; Roy, Jean-Francis; Yu, Bingqing; Chia, Patrick John; Bianchi, Federico; Cassani, Giovanni

*Published in:*  
arXiv

*Publication date:*  
2021

*Document Version*  
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

*Citation for published version (APA):*  
Tagliabue, J., Greco, C., Roy, J-F., Yu, B., Chia, P. J., Bianchi, F., & Cassani, G. (2021). SIGIR 2021 E-Commerce Workshop Data Challenge. *arXiv*.

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

### Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

# SIGIR 2021 E-Commerce Workshop Data Challenge

Jacopo Tagliabue  
Coveo Labs  
New York, USA  
jtagliabue@coveo.com

Ciro Greco  
Coveo Labs  
New York, USA  
cgreco@coveo.com

Jean-Francis Roy  
Coveo  
Quebec City, Canada  
jfroy@coveo.com

Bingqing Yu  
Coveo  
Montreal, Canada  
cyu2@coveo.com

Patrick John Chia  
Coveo  
Montreal, Canada  
pchia@coveo.com

Federico Bianchi  
Bocconi University  
Milano, Italy  
f.bianchi@unibocconi.it

Giovanni Cassani  
Tilburg University  
Tilburg, Netherlands  
g.cassani@tilburguniversity.edu

## ACM Reference Format:

Jacopo Tagliabue, Ciro Greco, Jean-Francis Roy, Bingqing Yu, Patrick John Chia, Federico Bianchi, and Giovanni Cassani. 2021. SIGIR 2021 E-Commerce Workshop Data Challenge. In *Montreal '21: SIGIR eCom, July, 2021, Montreal, CA*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 THE COVEO DATA CHALLENGE

The massive growth of e-commerce - which is estimated to turn into a 4.5 trillion industry in 2021 [37] - has fueled significant investments in areas such as Information Retrieval, NLP and recommendation systems. One of the principal goals of such investments has been to improve personalization in the customer journey. Optimal customer journeys in e-commerce are ideally supported by systems that transform the experience of each customer based on customer behavior and (a minimal set of) business rules [43].

Although compelling progress has been made, interesting challenges remain in personalization. In particular, customer intentions can be different depending on the occasion and even change within the same session, depending on where users are in their customer journey. In this context, the feedback loop determined by behavioural signals - such as searches, clicks, add-to-cart, cart-abandonment - spans from hours to a few seconds [38] and machine learning models need to adapt as fast as possible to the continuously changing nature of the customer journey.

The need for efficient procedures for personalization is even clearer if we consider the e-commerce landscape more broadly. Outside of giant digital retailers, the constraints of the problem are stricter, due to smaller user bases and the realization that most users are not frequently returning customers: in dozens of deployments across several verticals, we found that it is not uncommon

that only <30% of shoppers visit more than 2 times a year. Because of these constraints, the profit increase due to personalization strategies [10] has been often disappointing. In a context where user intentions are constantly changing, customers are mostly non-recurring and bounce rates are high [36], the competitive pressure towards real-time personalization requires machine learning models to work *as early as possible* in the customer journey and with *as little data* as possible. In this perspective, sessions emerge as the natural partition for building models with strong business impact [12, 23, 27, 39, 48] and session-based inference a crucial topic for retailers *of all sizes* [3, 41]

For practitioners reading about newer models, it is often hard to understand how those advances translate into their production context (if at all [26]), as researchers either used old (e.g. *MovieLens* [11] in [22, 27, 35]), industry-specific datasets (e.g. playlists [6]), or private datasets that may not be representative of common constraints [42]. For these reasons, we designed the **Coveo Data Challenge** with two principles in mind:

- advancing the field with the release of a new, fine-grained, session-based dataset. Our dataset (Section 3) presents product-level data from a mid-sized e-commerce, coupled with search-based interactions and content-based information; as such, it is a representative example of thousands of shops in the middle of the industry long tail, with realistic traffic, conversion rate, abandonment ratio;
- proposing inference tasks that are interesting for retailers of all sizes, but where models are evaluated on a mid-size shop. Given the increasing commoditization of tools and libraries, we expect more shops to develop ML capabilities in the coming years: as such, we believe in the value of research with tangible return on investment even outside of retail giants.

The rest of *this* document describes the data release, the chosen tasks and the evaluation procedure. Moreover, a brief survey of existing work is provided in Section 5.

### 1.1 About the organizers

The **Coveo Data Challenge** is organized by a collaboration between academic researchers and Coveo, a multi-tenant provider of

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Montreal '21, July, 2021, Montreal, CA

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

A.I. services, with a network of more than 1000 deployments for customer service, e-commerce and enterprise search use cases.

## 2 TASK DESCRIPTION

For this data challenge, we release a new session-level dataset, containing 10M product interactions over 57K different products (Section 3). Starting from this dataset, we accept solutions addressing two different tasks:

- (1) a **session-based recommendation task**, where a model is asked to predict the next interactions between shoppers and products, based on the previous product interactions and search queries within a session;
- (2) a **cart-abandonment task**, where, given a session containing an add-to-cart event for a product  $X$ , a model is asked to predict whether the shopper will buy  $X$  or not in that session. In particular, we are interested in modelling how abandonment prediction changes 1, 2, ...,  $n$  events after the original event, with the idea that the biggest business value comes from early prediction.

## 3 DATA DESCRIPTION

The **Coveo Data Challenge** features the release of a dataset comprising three main files for training:

- **browsing interactions**: a text file containing fine-grained shopping tracking for more than 4M sessions<sup>1</sup>. Each row is a separate interaction (page view or product interaction); events with the same `session_id_hash` are from the same session – please note that a product page may generate both a *detail* and a *pageview* event, and that the order of the events in the text file is not strictly chronological (refer to the session and timestamp information to reconstruct the actual chain of events for a given session). The fields are similar to the dataset in [33]<sup>2</sup>: `session_id_hash`, `event_type`, `product_action`<sup>3</sup>, `product_sku_hash`, `server_timestamp_epoch_ms`<sup>4</sup>, `hashed_url`. Descriptive statistics for this file can be found in Table 1;
- **search interactions**: a text file containing more than 800k search-based interactions. Each row is a search query event issued by a shopper, which includes an array of (hashed) results returned to the client. We also provide which result(s) have been clicked from the result set, if any. By also reporting products seen but not clicked, we hope to inspire clever ways to use negative feedback. The fields are: `session_id_hash`, `server_timestamp_epoch_ms`, `query_vector`, `product_skus_hash`, `clicked_skus_hash`. The field `query_vector` is a vectorized representation of the original search query, obtained through standard pre-trained models and dimensionality reduction techniques;

<sup>1</sup>The definition of “session” is the industry standard from <https://support.google.com/analytics/answer/2731565?hl=en>.

<sup>2</sup>The dataset is available at: <https://github.com/coveooss/shopper-intent-prediction-nature-2021>.

<sup>3</sup>Please note that removing a product from the cart might lead to multiple *remove* events, hence pre-processing might be needed with regard to this type of interactions.

<sup>4</sup>As a further anonymization technique, the timestamp has been shifted by an unspecified amount of weeks, keeping intact the intra-week patterns.

**Table 1: Descriptive statistics for the training dataset containing browsing interactions.**

Property	Value
# of SKUs	57,483
# of sessions	4,934,699
# of all events	36,079,307
# of all product interactions	10,431,611
# of add-to-cart	329,557
# of purchases	77,848
25/50/75 pct. of session length	2/3/8

- **product meta-data**: a text file containing a mapping between the SKUs (hashed) present in the entire dataset (training+testing) and their respective meta-data (when available), including a hashed representation of product categories, pricing information and vectorized representations of textual and image meta-data. The fields are: `product_sku_hash`, `category_hash`, `price_bucket`, `description_vector`, `image_vector`. The categories are hashed representations of a category tree<sup>5</sup>, where each level of hierarchy is separated with a `.`. The pricing information is provided as a 10-quantile integer. The vectorized fields are a dense representation of the original catalog, obtained through standard pre-trained modeling and dimensionality reduction techniques.

All files are available for download, and additional important information and sample scripts are filed under the project repository<sup>6</sup>. All training data comes from sampling in an unspecified ratio<sup>7</sup> several months of interactions in the life of *Shop Z*, a mid-size shop with *Alexa Ranking* between 25k and 200k; test data is sampled from a disjoint and adjacent time period. Data is fully anonymized and hashed, and collected in compliance with existing legislation: download and usage of the data implies acceptance of the terms and conditions (Section 6).

### 3.1 Comparison with existing datasets

Based on recent recommendation surveys and frameworks [24, 50] we can classify benchmarks in different categories including *ratings* (e.g. MovieLens [11], Anime [18], Netflix [29], Book-Crossing [51]), *reviews* (e.g. Epinions [49], Yelp [45], Steam [21], Amazon [28]), *advertising* (Criteo [17], Avazu [16]) and *e-commerce*. E-commerce datasets differ in terms of type of available data. Some only provide transactions (Ta-Feng, Belgium Retail Market, Foodarkt [31]) or product data for product classification [19, 20, 30], whereas others provide also interactions and product metadata (YOOCHOOSE (RCS15) [32], Retailrocket (RETAILR) [34], DIGINETICA [8]). Table 2 compares these three datasets with the data distributed for this challenge.

<sup>5</sup>The product categorization is given by the shop – in this case, product experts have defined a custom hierarchy to classify products for navigation and marketing purposes.

<sup>6</sup>The data challenge is hosted at <https://github.com/coveooss/SIGIR-ecom-data-challenge>.

<sup>7</sup>Exact traffic information is thus removed as a further anonymization strategy.

**Table 2: Comparison of e-commerce datasets from the literature with the Coveo dataset.**

	Coveo	DIGINETICA	RETAILR	RCS15
Searches	✓	✓		
Impressions	✓			
Clicks/Views	✓	✓	✓	✓
Carts	✓		✓	
Transactions	✓	✓	✓	✓
Category	tree	single	tree	—
Price	✓	✓	—	✓
Description	embedded	hashed name	hashed properties	—
Image	embedded	✓	—	—
Users	—	88K	1.4M	—
Items	66K	184K	235K	52K
Events	37M	3.3M	2.7M	42M

Considering the variety of the available features and the fine-grained details included in the **Coveo Data Challenge**, we believe it to be a significant addition to the set of resources in the space. In particular, by releasing neural-friendly embeddings for text, images *and* search queries, we hope to strike a virtuous trade-off between anonymity and information content.

## 4 EVALUATION

We describe here the evaluation procedure and the key metrics that will be used to evaluate submissions in our two tasks<sup>8</sup>. While the final ranking will solely reflect quantitative metrics, we encourage papers to include a section on qualitative evaluation, following the suggestions below.

### 4.1 Quantitative evaluation

- (1) **Recommendation.** We follow the evaluation protocol described in [24]: we reveal the start of a session to the model, and we ask it to predict future interactions. In particular, given a sequence of  $n$  events as input, we use two evaluation schemes:

- the model is evaluated at predicting the *immediate next item*, using the **Mean Reciprocal Rank** (MRR) as our metric;
- the model is evaluated at predicting *all subsequent items* of the session, up to a maximum of 20 after the current event. We take the **F1 score** from precision@20 and recall@20 as the final metric.

Moreover, we also measure additional quality factors defined in [14, 24]: coverage@20 and popularity bias@20. Finally, note that when a test session includes a search query, this event will be available to the model as an *input* (it can and should be taken into consideration by the model), but it will

<sup>8</sup>The evaluation script can be found in the public repository. Please refer to the challenge website for latest news about important dates, rules and submission format.

be ignored when computing the metrics: in other words, models are evaluated only on their ability to predict future interactions with products.

- (2) **Cart-abandonment.** Models will be evaluated by their **micro F1 score** in predicting shopping cart abandonment at the first add-to-cart event (AC), and then 2, 4, 6, 8 and 10 events after it.<sup>9</sup> Each model will be evaluated by performing a weighted combination of micro F1 scores at different clicks, with larger weights assigned to earlier predictions according to the following schema:

- micro F1 at AC \* 1
- micro F1 2 clicks after AC \* 0.9
- micro F1 4 clicks after AC \* 0.8
- micro F1 6 clicks after AC \* 0.7
- micro F1 8 clicks after AC \* 0.6
- micro F1 10 clicks after AC \* 0.5

The test set will not undergo any resampling to even the class distribution. Test sessions including purchases will be trimmed before the first purchase event to avoid trivializing the inference.

### 4.2 Qualitative evaluation

While we *do* recognize the importance of a standardized quantitative evaluation to have some measure of progress, we also agree with recent criticism of “leaderboard chasing” [9], as there are many practical constraints not captured by metrics which are crucial for the real-world success of a model.

To encourage a deeper understanding of the underlying business problems and foster a healthy competition among models with broad applicability, we solicit submissions that provide new insights by mixing quantitative results with industry-relevant qualitative discussion. As an inspiration, we provide below some examples of interesting topics, divided by task: as a general recommendation, reflections on training cost, inference time and sample efficiency are greatly appreciated, together with connections to related issues in IR for products [2].

- (1) **Recommendation.** We encourage papers to find solutions creatively using search and content data, together with browsing behavior. In particular, we encourage error analysis and comments on model performance for slices of the data that reflects practically important scenarios – new items, rare items, after-query prediction, interplay between search behavior and browsing patterns.
- (2) **Cart-abandonment.** We encourage submissions to include a careful error analysis which considers the relation between classification accuracy at different target clicks and different session characteristics, such as the number of interactions prior to the cart event and the total length of the session [33]. For example, suppose there are 15 actions after the AC in a purchase session: a model which converges on consistently predicting a purchase from action 11 (not captured in our evaluation) would anticipate the right outcome more than a model which is undecided until the second-to-last action

<sup>9</sup>By assessing performance at AC, moreover, we want to stress the importance of predicting intent from as little data as relevant.

but still correctly predicts a purchase at the end of the session. A measure of anticipation would be very informative although it depends on some free parameters (how strong should the support for a given prediction be, for how many consecutive actions should it remain stable, etc.): are there important insights of this sort that we can uncover leveraging this new dataset?

## 5 RELATED WORK

**Session-based models.** Session-based inference is a well-studied topic in e-commerce, as it powers recommendation [4, 7, 15, 44, 47] and NLP/IR use cases [40, 46]: while neural models are now dominating the research agenda [27, 39], there is also some skepticism about the widespread applicability of these techniques and the progress made by the field [25, 26]; in particular, a significant number of recent models still goes through sub-optimal benchmarks [22, 27, 35], which may not be representative of the type of constraints faced by many practitioners [13].

**Intent prediction.** The ability to predict user intent (add to cart, purchase, abandonment, etc.) from click-stream data is a long-standing challenge for e-commerce, with older solutions modelling user browsing in a Markovian fashion [1] and recent attempts leveraging deep learning architectures [5]. As for in-session recommendations, the broad impact of recent findings is unclear: for example, the architecture proposed in [42] was tested with an unrealistic conversion rate (10x a “standard” shop). More recently, [33] simulated accuracy in deployment scenarios by artificially sub-sampling raw data, showing that models outperform a simple baseline only up to a class imbalance of 80% to 20%. More research into models that scale to extreme class imbalance is thus a priority, and with our release we hope to nudge the community to devote more resources to this problem.

## 6 LEGAL NOTICE

The **Coveo Data Challenge** comes with specific Terms and Conditions, as found in the public repository and bundled with the dataset archive: by downloading and accessing the data, you agree to comply with the terms – in addition, by submitting to the challenge, you agree to release your code under an open source license. In particular, the dataset is freely provided to the research community with the purpose of advancing the field: all data was collected and processed in an anonymized fashion, in compliance with existing legislation, and released with additional hashing of meta-data: participants agree to not use the dataset for any other purpose than what is stated in the Terms and Conditions, nor attempt to reverse engineer or de-anonymise the dataset by explicitly or implicitly linking the data to any person, brand or legal entity.

## ACKNOWLEDGMENTS

The organizers wish to thank Luca Bigon for his outstanding support in data collection, and Surya Kallumadi, Massimo Quadrana, Dietmar Jannach, Ajinkya Kale for precious feedback on a previous version of this paper. Finally, special thanks to Richard Tessier and Coveo’s legal team for believing in this data sharing initiative.

## REFERENCES

- [1] Dimitris J. Bertsimas, Adam J. Mersereau, and Nitin R. Patel. [n.d.]. *Dynamic Classification of Online Customers*. 107–118. <https://doi.org/10.1137/1.9781611972733.10> arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9781611972733.10>
- [2] Federico Bianchi, Jacopo Tagliabue, and Bingqing Yu. 2021. Query2Prod2Vec: Grounded Word Embeddings for eCommerce. In *NAACL-HLT*. Association for Computational Linguistics.
- [3] Federico Bianchi, Jacopo Tagliabue, Bingqing Yu, Luca Bigon, and Ciro Greco. 2020. Fantastic Embeddings and How to Align Them: Zero-Shot Inference in a Multi-Shop Scenario. In *Proceedings of the SIGIR 2020 eCom workshop, July 2020, Virtual Event*. <https://arxiv.org/abs/2007.14906>
- [4] Federico Bianchi, Bingqing Yu, and Jacopo Tagliabue. 2020. BERT Goes Shopping: Comparing Distributional Models for Product Representations. *arXiv preprint arXiv:2012.09807* (2020).
- [5] Luca Bigon, Giovanni Cassani, Ciro Greco, L. Lacasa, Mattia Pavoni, Andrea Polonioli, and Jacopo Tagliabue. 2019. Prediction is very hard, especially about conversion. Predicting user purchases from clickstream data in fashion e-commerce. *ArXiv abs/1907.00400* (2019).
- [6] Hugo Caselles-Dupré, Florian Lesaint, and Jimena Royo-Letelier. 2018. Word2vec Applied to Recommendation: Hyperparameters Matter. In *Proceedings of the 12th ACM Conference on Recommender Systems (Vancouver, British Columbia, Canada) (RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 352–356. <https://doi.org/10.1145/3240323.3240377>
- [7] Wanyu Chen, Fei Cai, Honghui Chen, and Maarten de Rijke. 2019. A Dynamic Co-Attention Network for Session-Based Recommendation. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (Beijing, China) (CIKM '19)*. Association for Computing Machinery, New York, NY, USA, 1461–1470. <https://doi.org/10.1145/3357384.3357964>
- [8] CIKM-Cup. 2016. Cikum cup 2016 track 2: Personalized e-commerce search challenge. <https://competitions.codalab.org/competitions/11161>
- [9] Kawin Ethayarajh and Dan Jurafsky. 2020. Utility Is in the Eye of the User: A Critique of NLP Leaderboard Design. *ArXiv abs/2009.13888* (2020).
- [10] Penny Gillespie, Jason Daigler, Mike Lowndes, Christina Klock, Yanna Dharmasithira, and Sandy Shen. 2018. *Magic Quadrant for Digital Commerce*. Technical Report. Gartner.
- [11] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. <https://doi.org/10.1145/2827872>
- [12] Balázs Hidasi, Alexandros Karatzoglou, L. Baltrunas, and D. Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. *CoRR abs/1511.06939* (2016).
- [13] Dietmar Jannach, Gabriel de Souza P. Moreira, and Even Oldridge. 2020. Why Are Deep Learning Models Not Consistently Winning Recommender Systems Competitions Yet? A Position Paper. In *Proceedings of the Recommender Systems Challenge 2020 (Virtual Event, Brazil) (RecSysChallenge '20)*. Association for Computing Machinery, New York, NY, USA, 44–49. <https://doi.org/10.1145/3415959.3416001>
- [14] Dietmar Jannach and Malte Ludewig. 2017. When recurrent neural networks meet the neighborhood for session-based recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. 306–310.
- [15] Yu jia Zheng, Siyi Liu, Zekun Li, and S. Wu. 2020. DGTN: Dual-channel Graph Transition Network for Session-based Recommendation. *ArXiv abs/2009.10002* (2020).
- [16] Kaggle. 2014a. Avazu click-through rate prediction. <https://www.kaggle.com/c/avazu-ctr-prediction/data>
- [17] Kaggle. 2014b. Display advertising challenge. <https://www.kaggle.com/c/criteo-display-ad-challenge/data>
- [18] Kaggle. 2016a. Anime Recommendations Database. <https://www.kaggle.com/CooperUnion/anime-recommendations-database>
- [19] Kaggle. 2016b. eCommerce Item Data. <https://www.kaggle.com/cclark/product-item-data/home>
- [20] Kaggle. 2017. Innerwear Data from Victoria’s Secret and Others. <https://www.kaggle.com/PromptCloudHQ/innerwear-data-from-victorias-secret-and-others>
- [21] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [22] Walid Krichene and Steffen Rendle. 2020. On sampled metrics for item recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1748–1757.
- [23] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: Short-Term Attention/Memory Priority Model for Session-based Recommendation. 1831–1839. <https://doi.org/10.1145/3219819.3219950>
- [24] Malte Ludewig and Dietmar Jannach. 2018. Evaluation of session-based recommendation algorithms. *User Modeling and User-Adapted Interaction* 28, 4-5 (2018), 331–390.
- [25] Malte Ludewig, Noemi Mauro, Sara Latifi, and D. Jannach. 2019. Empirical Analysis of Session-Based Recommendation Algorithms. *ArXiv abs/1910.12781*

- (2019).
- [26] Malte Ludewig, Noemi Mauro, Sara Latifi, and Dietmar Jannach. 2019. Performance Comparison of Neural and Non-Neural Approaches to Session-Based Recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (*RecSys '19*). Association for Computing Machinery, New York, NY, USA, 462–466. <https://doi.org/10.1145/3298689.3347041>
- [27] Yifei Ma, Balakrishnan (Murali) Narayanaswamy, Haibin Lin, and Hao Ding. 2020. Temporal-Contextual Recommendation in Real-Time. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) (*KDD '20*). Association for Computing Machinery, New York, NY, USA, 2291–2299. <https://doi.org/10.1145/3394486.3403278>
- [28] 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*. 43–52.
- [29] Netflix. 2009. Netflix Prize Data. <https://netflixprize.com/>
- [30] Rakuten. 2020. SIGIR eCom'20 Data Challenge: Multimodal Product Classification and Retrieval. <https://sigir-ecom.github.io/data-task.html>
- [31] RecSys. 2015. Grocery Shopping Datasets. [https://recsyswiki.com/wiki/Grocery\\_shopping\\_datasets](https://recsyswiki.com/wiki/Grocery_shopping_datasets)
- [32] Recsys-Challenge. 2015. Recsys challenge 2015. <https://2015.recsyschallenge.com/challenge.html>
- [33] Borja Requena, Giovanni Cassani, Jacopo Tagliabue, Ciro Greco, and Lucas Lacasa. 2020. Shopper intent prediction from clickstream e-commerce data with minimal browsing information. *Scientific Reports* 10 (2020), 2045–2322. <https://doi.org/10.1038/s41598-020-73622-y>
- [34] Retailrocket. 2016. Retailrocket recommender system dataset. <https://www.kaggle.com/retailrocket/ecommerce-dataset>
- [35] Tobias Schnabel and Paul N Bennett. 2020. Debiasing Item-to-Item Recommendations With Small Annotated Datasets. In *Fourteenth ACM Conference on Recommender Systems*. 73–81.
- [36] SimilarWeb. 2019. *Top sites ranking for E-commerce And Shopping in the world*. Retrieved December 23, 2019 from <https://www.similarweb.com/top-websites/category/e-commerce-and-shopping>
- [37] Statista Research Department. 2020. *Global retail e-commerce sales 2014-2023*. Retrieved November 29, 2020 from <https://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/>
- [38] Ning Su, Jiyin He, Yiqun Liu, Min Zhang, and Shaoping Ma. 2018. User Intent, Behaviour, and Perceived Satisfaction in Product Search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (*WSDM '18*). Association for Computing Machinery, New York, NY, USA, 547–555. <https://doi.org/10.1145/3159652.3159714>
- [39] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (*CIKM '19*). Association for Computing Machinery, New York, NY, USA, 1441–1450. <https://doi.org/10.1145/3357384.3357895>
- [40] Jacopo Tagliabue and Bingqing Yu. 2020. Shopping in the Multiverse: A Counterfactual Approach to In-Session Attribution. In *Proceedings of the SIGIR 2020 eCom workshop, July 2020, Virtual Event*.
- [41] Jacopo Tagliabue, Bingqing Yu, and Marie Beaulieu. 2020. How to Grow a (Product) Tree. Personalized Category Suggestions for eCommerce Type-Ahead. In *Companion Proceedings of ACL* (Seattle, USA). Association for Computing Machinery, New York, NY, USA.
- [42] Arthur Toth, L. Tan, G. Fabbri, and Ankur Datta. 2017. Predicting Shopping Behavior with Mixture of RNNs. In *eCOM@SIGIR*.
- [43] Manos Tsagkias, Tracy Holloway King, Surya Kallumadi, Vanessa Murdock, and Maarten de Rijke. 2020. Challenges and Research Opportunities in eCommerce Search and Recommendations. In *SIGIR Forum*, Vol. 54.
- [44] S. Wu, Y. Tang, Yanqiao Zhu, L. Wang, X. Xie, and T. Tan. 2019. Session-based Recommendation with Graph Neural Networks. In *AAAI*.
- [45] Yelp. 2020. Yelp dataset. <https://www.yelp.com/dataset>
- [46] Bingqing Yu, Jacopo Tagliabue, Ciro Greco, and Federico Bianchi. 2020. “An Image is Worth a Thousand Features”: Scalable Product Representations for In-Session Type-Ahead Personalization. In *Companion Proceedings of the Web Conference 2020* (Taipei, Taiwan) (*WWW '20*). Association for Computing Machinery, New York, NY, USA, 461–470. <https://doi.org/10.1145/3366424.3386198>
- [47] Feng Yu, Yanqiao Zhu, Qiang Liu, S. Wu, L. Wang, and Tieniu Tan. 2020. TAGNN: Target Attentive Graph Neural Networks for Session-based Recommendation. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2020).
- [48] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M. Jose, and Xiangnan He. 2019. A Simple Convolutional Generative Network for Next Item Recommendation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Melbourne VIC, Australia) (*WSDM '19*). Association for Computing Machinery, New York, NY, USA, 582–590. <https://doi.org/10.1145/3289600.3290975>
- [49] Tong Zhao, Julian McAuley, and Irwin King. 2014. Leveraging social connections to improve personalized ranking for collaborative filtering. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*. 261–270.
- [50] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Kaiyuan Li, Yushuo Chen, Yujie Lu, Hui Wang, Changxin Tian, Xingyu Pan, et al. 2020. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. *arXiv preprint arXiv:2011.01731* (2020).
- [51] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*. 22–32.