

---

# GreenDB - A Dataset and Benchmark for Extraction of Sustainability Information of Consumer Goods

---

Alexander Flick<sup>\*1</sup> Sebastian Jäger<sup>\*1</sup> Jessica Adriana Sanchez Garcia<sup>1</sup> Kaspar von den Driesch<sup>1</sup>  
Karl Brendel<sup>1</sup> Felix Biessmann<sup>1,2</sup>

## Abstract

The production, shipping, usage, and disposal of consumer goods have a substantial impact on greenhouse gas emissions and the depletion of resources. Machine Learning (ML) can help to foster sustainable consumption patterns by accounting for sustainability aspects in product search or recommendations of modern retail platforms. However, the lack of large high quality publicly available product data with trustworthy sustainability information impedes the development of ML technology that can help to reach our sustainability goals. Here we present *GreenDB*, a database that collects products from European online shops on a weekly basis. As proxy for the products' sustainability, it relies on sustainability labels, which are evaluated by experts. The *GreenDB* schema extends the well-known *schema.org Product* definition and can be readily integrated into existing product catalogs. We present initial results demonstrating that ML models trained with our data can reliably (F1 score 96%) predict the sustainability label of products. These contributions can help to complement existing e-commerce experiences and ultimately encourage users to more sustainable consumption patterns.

2021a) and the depletion of resources. With every purchase decision made, people can decide for – or against – more sustainable consumption behavior. While many consumers would like to choose sustainable options (PwC, 2019; Economist Intelligence Unit, 2021; Umweltbundesamt, 2021b) at the point of sale, the relevant information is not available or difficult to evaluate for the consumers due to lacking standards. There are attempts to take sustainability information and user preferences into account for recommender systems (Tomkins et al., 2018). Our research and discussions with researchers, managers, and engineers working for large online retailers have shown that the most important factor hindering the integration of sustainability aspects in their shopping experiences is the availability of trustworthy sustainability data. To foster more sustainable shopping behavior, we need to build ML methods that are aware of sustainability. This requires large high-quality datasets that are recent, structured, and trustworthy<sup>1</sup>. To the best of our knowledge, the largest publicly available database containing products and their respective sustainability information is the *GreenDB* (Jäger et al., 2022). As proxy for the products' sustainability, it relies on sustainability labels, which are evaluated by experts regarding their *credibility*, *environmental*, and *socio-economic* aspects. Currently, the *GreenDB* includes about 230,000 unique products from 4 shops in 2 European countries. However, it is growing continuously and is updated at a weekly cadence. Its 26 product categories are selected based on a careful analysis of search logs of several millions of users of one of the largest search engines in Europe. An overview of the data in the *GreenDB* is shown in Figure 1. For detailed information about the *GreenDB* data collection system and sustainability evaluation process, we refer the reader to (Jäger et al., 2022).

## 1. Introduction

Climate change is one of the most important challenges of our generation. The increasing demand for consumer goods and their implications for globalized supply chains, shipping, and disposal are major factors in increased greenhouse gas emissions (Ritchie & Roser, 2020; Umweltbundesamt,

## 2. Related Work

To better assess the value and contributions made with the *GreenDB*, we present an overview of related public

---

<sup>\*</sup>Equal contribution <sup>1</sup>Berliner Hochschule für Technik <sup>2</sup>Einstein Center Digital Future, Berlin. Correspondence to: Sebastian Jäger <sebastian.jaeger@bht-berlin.de>.

---

<sup>1</sup>Many retailers and manufacturers provide sustainability information that is primarily intended for marketing purposes (greenwashing).

datasets that complement our work. The Web Data Commons (WDC) (Primpeli et al., 2019) leverages structured metadata from websites HTML to create product datasets for product matching and product categorization. Based on the product corpus, they have published the WDC-25 Gold Standard for Product Categorization, consisting of 24,000 product offers from different e-shops, which were manually assigned to 25 product categories. The Amazon ML Challenge aims at amazon browse node (product category) classification based on product metadata like *name*, *description*, and *brand*. The dataset consists of about 2.9 million products, each assigned to one of 9,919 browse nodes. (Chalenge, 2022) The Amazon Review Data focuses on product reviews. However, it also contains product attributes such as *color* or *size* and offers many product categories. Since the data spans from 1996 to 2018, the dataset size is enormous (about 34GB). (Ni et al., 2019) The Klarna Product-Page Dataset was collected between 2018 and 2019 and is intended to benchmark representation learning algorithms on the task of web element prediction on e-commerce websites. It consists of 51,701 product pages as HTML and screenshots of the website and has been manually labeled, specifically, the product *price*, *name*, *image*, *add-to-cart*, and *go-to-cart buttons*. In addition, the data comes from 8 different markets (US, GB, SE, NL, FI, NO, DE, AT). (Hotti et al., 2021) Flipkart Products is most similar to GreenDB. It is a tabular dataset with 20,000 products and 15 columns. These attributes are partially very similar to ours, e.g., *product\_name*, *product\_url*, *retail\_price*, *description*, and *brand*. (Flipkart, 2022) However, none of them integrates products’ sustainability information, nor are they updated regularly. More recently, sustainability certification institutions aim at publishing their data, such as the *European Product Database for Energy Labeling (EPREL)* (European Union, 2022), *Blue Angel* (Ger, 2022), or *Higg Sustainability Profiles* (Hig, 2022). These data sets are rather small and contain very specific selections of products.

### 3. Predicting Sustainability Information

To establish a simple baseline for automated inference of sustainability information from online products, we used the *TabularPredictor* module of the AutoML package *AutoGluon* (AutoGluon; Erickson et al., 2020). The model was trained on the entire GreenDB dataset (Jäger et al., 2022). We simplified the prediction problem, without loss of generality, from a multilabel problem (each product can have more than one sustainability label) to a multiclass prediction problem: For each product, we consider the first sustainability label (sorted in alphabetical order). In Table 1, we list the results. Note that our simple benchmark achieves an F1 score (class frequency weighted) of 0.96, indicating that the high data quality of the GreenDB enables researchers

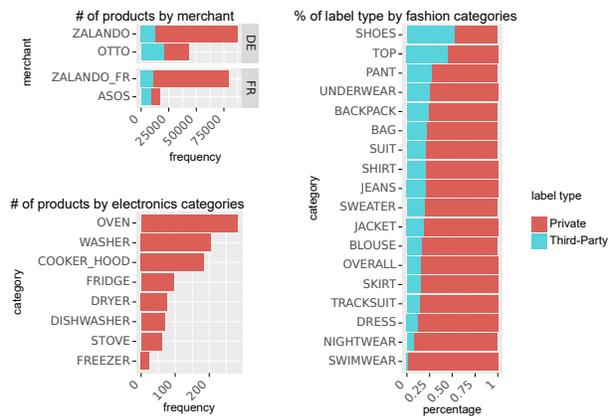


Figure 1. GreenDB overview. We follow the terminology by (Gossen et al., 2022) and differentiate between *private* and *third-party* labels. Third-party labels require independent verification of the label’s environmental and/or socio-economic requirements. However, private labels are self-declared claims and not based on independent verification. Depending on the online shops and product categories, the number (and fraction) of preferable third-party labels varies and is generally much lower.

	precision	recall	f1
BETTER_COTTON_INITIATIVE	0.89	0.93	0.91
BIORE	0.98	0.98	0.98
BLUESIGN_APPROVED	0.92	0.92	0.92
BLUESIGN_PRODUCT	0.95	0.97	0.96
COTTON_MADE_IN_AFRICA	0.95	0.98	0.97
CRADLE_TO_CRADLE_BRONZE	0.33	0.14	0.20
CRADLE_TO_CRADLE_GOLD	0.54	0.24	0.33
CRADLE_TO_CRADLE_SILVER	0.75	0.27	0.40
ECOCERT	0.94	0.68	0.79
EU_ECOCOLABEL_TEXTILES	1.00	0.91	0.95
FAIRTRADE_COTTON	0.97	0.85	0.90
FAIRTRADE_TEXTILE_PRODUCTION	1.00	0.80	0.89
GLOBAL_RECYCLED_STANDARD	0.93	0.85	0.89
GOOD_CASHMERE_STANDARD	0.00	0.00	0.00
GOTS_MADE_WITH_ORGANIC_MATERIALS	0.84	0.40	0.54
GOTS_ORGANIC	0.85	0.79	0.82
GREEN_BUTTON	0.96	0.88	0.92
LEATHER_WORKING_GROUP	0.93	0.93	0.93
MADE_IN_GREEN_OEKO_TEX	0.98	0.95	0.96
ORGANIC_CONTENT_STANDARD_100	0.79	0.59	0.67
ORGANIC_CONTENT_STANDARD_BLENDED	0.62	0.41	0.49
PRIVATE_LABEL	0.97	0.98	0.98
RECYCLED_CLAIM_STANDARD_100	0.77	0.50	0.61
RECYCLED_CLAIM_STANDARD_BLENDED	1.00	0.93	0.97
RESPONSIBLE_DOWN_STANDARD	0.94	0.72	0.81
RESPONSIBLE_WOOL_STANDARD	0.72	0.54	0.62
accuracy	0.96	0.96	0.96
macro avg	0.83	0.69	0.74
weighted avg	0.96	0.96	0.96

Table 1. Prediction results of sustainability labels based on the GreenDB product attributes *name*, *description*, *brand*, *category*, *color*, *merchant*, and *price*. The extraction of these attributes is part of the GreenDB data collection system (Jäger et al., 2022).

to quickly achieve production-ready ML-based automated sustainability inference.

## References

- The German Ecolabel, 2022. URL <https://www.blauer-engel.de/en>. [Online; accessed 23-May-2022].
- Higg Sustainability Profiles, 2022. URL <https://profiles.higg.com/>. [Online; accessed 23-May-2022].
- AutoGluon. <https://github.com/awsml/autogluon>.
- Challenge, A. M., 2022. URL <https://www.hackerearth.com/en-us/challenges/competitive/amazon-ml-challenge/>. [Online; accessed 23-May-2022].
- Economist Intelligence Unit. An eco-wakening - measuring global awareness, engagement and action for nature, 2021. URL [https://files.worldwildlife.org/wfcmprod/files/Publication/file/93ts5bhvyq\\_An\\_EcoWakening\\_Measuring\\_awareness\\_engagement\\_and\\_action\\_for\\_nature\\_FINAL\\_MAY\\_2021.pdf](https://files.worldwildlife.org/wfcmprod/files/Publication/file/93ts5bhvyq_An_EcoWakening_Measuring_awareness_engagement_and_action_for_nature_FINAL_MAY_2021.pdf).
- Erickson, N., Mueller, J., Shirkov, A., Zhang, H., Larroy, P., Li, M., and Smola, A. Autogluon-tabular: Robust and accurate autogl for structured data. *arXiv preprint arXiv:2003.06505*, 2020.
- European Union. Product database, 2022. URL [https://ec.europa.eu/info/energy-climate-change-environment/standards-tools-and-labels/products-labelling-rules-and-requirements/energy-label-and-ecodesign/product-database\\_de](https://ec.europa.eu/info/energy-climate-change-environment/standards-tools-and-labels/products-labelling-rules-and-requirements/energy-label-and-ecodesign/product-database_de). [Online; accessed May 19, 2022].
- Flipkart, 2022. URL <https://www.kaggle.com/PromptCloudHQ/flipkart-products>. [Online; accessed 23-May-2022].
- Gossen, M., Jäger, S., Hoffmann, M. L., Bießmann, F., Korenke, R., and Santarius, T. Nudging Sustainable Consumption: A Large-Scale Data Analysis of Sustainability Labels for Fashion in German Online Retail. *Frontiers in Sustainability*, 3:922984, June 2022. ISSN 2673-4524. doi: 10.3389/frsus.2022.922984. URL <https://www.frontiersin.org/articles/10.3389/frsus.2022.922984/full>.
- Hotti, A., Risuleo, R. S., Magureanu, S., Moradi, A., and Lagergren, J. The klarna product page dataset: A realistic benchmark for web representation learning. *CoRR*, abs/2111.02168, 2021. URL <https://arxiv.org/abs/2111.02168>.
- Jäger, S., Greene, J., Jakob, M., Korenke, R., Santarius, T., and Biessmann, F. GreenDB: Toward a Product-by-Product Sustainability Database. Technical Report arXiv:2205.02908, arXiv, May 2022. URL <http://arxiv.org/abs/2205.02908>.
- Jäger, S., Bießmann, F., Flick, A., Sanchez Garcia, J. A., von den Driesch, K., and Brendel, K. GreenDB: A Product-by-Product Sustainability Database, February 2022. URL <https://doi.org/10.5281/zenodo.6576662>. Supported by the Federal Ministry for the Environment, Nature Conservation and Nuclear Safety based on a decision of the German Bundestag. Förderkennzeichen: 67KI2022B.
- Ni, J., Li, J., and McAuley, J. J. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 188–197. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1018. URL <https://doi.org/10.18653/v1/D19-1018>.
- Primpeli, A., Peeters, R., and Bizer, C. The wdc training dataset and gold standard for large-scale product matching. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, pp. 381–386, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450366755. doi: 10.1145/3308560.3316609.
- PwC. Surviving the retail apocalypse, 2019. URL <https://www.pwc.de/de/handel-und-konsumguter/studie-surviving-the-retail-apocalypse.pdf>.
- Ritchie, H. and Roser, M. Co2 and greenhouse gas emissions. *Our World in Data*, 2020. <https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions>.
- Tomkins, S., Isley, S., London, B., and Getoor, L. Sustainability at scale: towards bridging the intention-behavior gap with sustainable recommendations. In Pera, S., Ekstrand, M. D., Amatriain, X., and O'Donovan, J. (eds.), *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, pp. 214–218. ACM, 2018. doi: 10.1145/3240323.3240411. URL <https://doi.org/10.1145/3240323.3240411>.
- Umweltbundesamt. Konsum und umwelt: Zentrale handlungsfelder, 2021a. URL <https://www.umweltbundesamt.de/themen/wirtschaft-konsum/konsum-umwelt-zentrale-handlungsfelder>. [Online; accessed February 03, 2022].
- Umweltbundesamt. Umweltbewusstsein in deutschland, 2021b. URL <https://www.umweltbundesamt.de/themen/nachhaltigkeit-strategien-internationales/umweltbewusstsein-in-deutschland>. [Online; accessed February 03, 2022].