Product Reviews as Source for Extracting Product Information – Lessons Learned

Kunal Kunal[1,*], Norbert Fuhr[1] and Andrea Papenmeier[2]

[1]*University of Duisburg-Essen, Forsthausweg 2, 47057 Duisburg, Germany*
[2]*University of Twente, Drienerlolaan 5, 7522 NB Enschede, Netherlands*

Abstract

E-commerce platforms offer access to rich information about products, including customer-written reviews. This new data is used as signal for result list ranking, dynamic facets, or to generate product descriptions. In this paper, we put review data to the test and analyze their potential to serve as a reliable data source for information about products. For 50 products, we compared product details mentioned in reviews with the actual product information and identified several pitfalls, including customers talking about different products, reporting context-dependent values, and uttering about desired product specifications. With this work, we highlight challenges that need to be accounted for when employing automatically extracting information from customer review in e-commerce.

Keywords
Information Extraction, E-Commerce, Review Data, Product Description

1. Introduction

Customers of e-commerce platforms often rely on product reviews to assess a product and the potential risks of purchasing it [1], e.g., product quality, shipping times, or customer service issues. Ultimately, reviews influence customers' purchase decisions – whether a customer decides to buy a product or choose one brand over another [2, 3]. Additionally, reviews can highlight important product attributes that may not be evident in seller-provided descriptions [4]. On the other side, reviews provide valuable feedback to sellers: revealing what customers like or dislike about an item allows businesses to improve the products [5].

While the economic relevance of online shopping is increasing [6], so is the volume of online reviews [7]. Customers find it increasingly difficult to distinguish between valuable and worthless reviews due to the sheer volume of available reviews and their inconsistent quality, which reduces the usefulness of online reviews [8]. As the number of reviews start to exceed the human cognitive processing capacity, e-commerce platforms started to adopt automated tools in order to keep the benefits of reviews [9]. Previous works have proposed to extract product information from reviews with automated methods, e.g., to provide a comprehensive overview of other customers' options [9] or to enable customers to filter for reviews talking

about a specific product aspect [10]. Moreover, information extracted from reviews is used to adjust the ranking of products [11, 12, 13], to build dynamic facets for filtering [14], or identify user needs for the design of future products [15]. Reviews have also been used to generate informative product descriptions that deliver more information about the product than badly-written product specifications from the seller [16].

However, using automated methods to extract information about the products from its reviews requires reliable and truthful data. Prior research, so far, has often evaluated automated methods quantitatively with only few diving into the reasons for incorrect instances. Most identified challenges relate to NLP problems such as the use of synonyms, abbreviations, and variations in writing styles [17, 18]. Some earlier works also mention semantic challenges such as reviews mentioning contextual information that are falsely recognized as product information [19] or irrelevant subjects being talked about [17]. To what extent those observations still hold true nowadays remains to be investigated.

We therefore set out to analyse the quality of review texts in e-commerce to serve as data source for reliable information extraction. In this research, we investigated the use case of laptop products and analyzed 1500 review texts spread over 50 laptop products. We manually annotated and compared information from reviews with the listed product information and clustered discrepancies into five general and one attribute-dependent "lessons learned". We focused on eight product aspects for which the product description – when complete – serves as ground truth[1].

Our investigation uncovered several challenges for automated information extraction from product reviews, such as mentions of information about different products, neglect of formatting standards, and reporting of approximate values and experienced real-life values that do not match with the seller's statements about maximum values. The findings highlight the necessity for methodologies and approaches to effectively identify and extract product specification problems amidst the myriad of customer reviews.

2. Method

In this work, we set out to conduct an exploratory study for investigating the potential of review data for extracting information about product aspects, guided by the following research questions:

RQ1: *To what extent do review texts contain information about product aspects in e-commerce?*

RQ2: *Are there discrepancies between information mentioned in review texts and the actual specifications of products?*

For this initial study, we choose a product domain that is commonly known and used, leading to a wide body of reviews per product. Additionally, to answer the research questions, ground truth data of product characteristics is needed. We therefore chose the use case of laptop products, with their complex, multi-faceted aspects, and large availability of review data.

---

[1]As less than 20% of product descriptions are actually complete (cf. Section 3.1), one could also hope to extract the missing information from the reviews.

Laptops are a type of utilitarian product with available ground truth data and a wide range of (technical) product aspects that are discussed in the reviews. In our research we focus on eight product aspects: processor, RAM, screen, battery, model, brand, price, and storage. We manually annotated laptop reviews regarding utterances about the chosen product aspects, and compared them with the seller-generated product information. We then identified cases with discrepancies and clustered the causes for discrepancies across product aspects.

2.1. Product Dataset with Reviews

We first collected a product dataset of about 5000 laptops including reviews from a prominent e-commerce platform. We eliminated products that were missing seller-generated information about the product aspects relevant for this experiment (processor, RAM, screen, battery, model, brand, price, storage). Of the remaining 805 products, we chose products with at least 30 reviews, and drew a random subset of 50 products from the remaining set. We kept the 30 longest reviews for each product, assuming that the longest reviews contain most information about the product. The final dataset for analysis contains 1500 review texts: 50 products with 30 reviews each. Listing 1 displays an excerpt of a product data point.

Listing 1: Example of a product object in the dataset

```
1  {
2    "title": "Lenovo X1 Carbon 6th Generation Ultrabook: Core i7-8550U, 16GB RAM, 512GB
          SSD, 14inch Full HD Display, Backlit Keyboard (Renewed)",
3    "brand": "Lenovo",
4    "model": "X1 Carbon",
5    "price": 330.98,
6    "batteryLife": "15.0 hours",
7    "totalStorageCapacity": "512 gigabytes",
8    "storageType": "SSD",
9    "processorCores": "4-core",
10   "processorBrand": "Intel",
11   "processorSpeedBase": "1.8 gigahertz",
12   "systemMemoryRam": "16 gigabytes",
13   "systemMemoryRamType": "DDR3 SDRAM",
14   "screenSize": "14.0 inches",
15   "screenResolution": "1920 x 1080",
16   "screenResolutionName": "Full HD",
17   ...
18 }
```

2.2. Annotations

To investigate the *informational content* in review texts for laptop products (**RQ1**), we annotated relevant utterances from all 1500 laptop product reviews pertaining to the product aspects under investigation (processor, RAM, screen, battery, model, brand, price, storage). The utterances spanned one to multiple words and contained either precise, technical information (e.g., *"256GB"* for storage, *"HP"* for brand) or vague, natural language descriptions of an aspect (e.g., *"long battery life"* for battery, *"the screen will not turn on"* for screen).

2.3. Mock Extraction

To investigate the *correctness* of the information mentioned in reviews (**RQ2**), we simulated the task of automatically extracting precise, technical information about a product. Previous research often focused on extracting sentiments and opinions from reviews [20, 13], with some works also accounting for the technical specifications of a product [21]. In this study, we focused on technical information because it can be compared to the seller-generated product information which serves as ground truth. For the eight annotated product aspects[22], we aimed to determine whether their precise, technical values could be determined from the product's reviews. From the annotations of a laptop's reviews, In this research we manually extracted all precise, technical values per aspect[23] and used majority voting over this set of values. In this study, we aim to perform a manual annotation to identify all existing issues comprehensively without adding mistakes from automatic methods. Subsequent research could address the model extraction problem automatically with large language models, such as the Llama model, and propose potential solutions.

For example, for one of the laptops in the dataset, we found the following seven utterances regarding storage in its 30 reviews:

> *"not satisfied with the memory space", "64GB SSD", "32GB SSD", "storage is minimal",*
> *"had to buy extra storage", "32 gig", "64 gig"*

This example leads to a tie situation because there are two occurrences of 64GB and two occurrences of 32GB.

Additionally, for each case where the value in the review utterance differed from the actual value of the product, identified the reason for the discrepancy and conducted two cycles of thematic analysis [24].

3. Results

To answer **RQ1**, we annotated utterances with aspect information in review texts. For example, the following review text contains four annotated utterances (in bold; one about the storage, three about the screen):

> *"The **solid state hard drive makes no noise**. The **full HD screen** is **very crisp***
> *with **vibrant colors**."*

Table 1 shows how many products have reviews with information about each of the eight aspects. For battery, brand, and price, we could find information about the respective aspect in the reviews of all products. RAM was the least popular aspect in reviews: Only 29 products had mentions of RAM sizes in their reviews. For the use case of extracting technical information, reviews seem to provide some data, but not for all products. Brand names and battery life values in hours were popularly talked about in reviews. However, RAM sizes, model names, and screen sizes were not often mentioned in reviews. For those aspects, review texts do not provide a lot of data for automated extraction.

Overall, we found many utterances in the reviews that contained information about specific laptop aspects. In a second step, we also investigated the *correctness* of that information since we need reliable and correct information in the data that we use in automated extraction (**RQ2**).

Table 1
Product counts with utterances of precise or non-precise values of aspects.

Products with...	Processor	RAM	Screen	Battery	Model	Brand	Price	Storage
No Mentions of Aspect	9	21	4	0	25	0	0	4
Mentions of Aspect	41	29	46	50	25	50	50	46
∟ **Precise Values**	**33**	**17**	**25**	**41**	**21**	**50**	**42**	**31**
∟ Only Vague Values	8	12	21	9	4	0	8	15

Table 2 reports the results of the mock extraction analysis in which we manually extracted precise product information from the annotated review texts. The findings show a discrepancy between product information stated by the seller and product information mentioned in customer reviews: For many products, the mock extraction determined incorrect values. For example, for laptop screens, 25 out of 50 products have mentions of precise screen sizes or resolutions in their reviews. Only in 4 cases, the correct value would have been determined using majority voting. In 13 cases, an incorrect value would have been determined. That is, the reviews of a product contain more often an incorrect value (e.g., *"15 inches big"* instead of 15.6 inches). In 8 cases, we cannot determine a value because multiple possible values have the same occurrence, leading to a tie.

Table 2
Results of mock value extraction with majority voting for precise mentions.

Products with...	Processor	RAM	Screen	Battery	Model	Brand	Price	Storage
(i) No Values for Extraction	17	33	25	9	29	0	8	19
(ii) Correct Value Extracted	17	9	4	0	16	37	0	18
(iii) Incorrect Value Extracted	16	1	13	41	5	13	42	4
(iv) Tie	0	7	8	0	0	0	0	9

4. Lessons Learned

The mock value extraction shows that automatically extracting information from reviews can lead to errors. Using thematic analysis, we identified five general causes for discrepancies. To elaborate on **RQ2**, we propose the following five lessons learned for using reviews as data source for extracting product information in e-commerce.

4.1. Lesson 1: Cross-Product Comparisons

In some cases, customers referred to other products in their reviews.

Previous Products Customers mentioned their previous laptop across all product aspects. They frequently reference their experiences with previous products, highlighting specific values associated with these products to contextualize their evaluations. For instance, when discussing the processor, customers often mention the model and specifications of their old processors and compared them to the processor of the reviewed laptop:

> *"This CORE i7 8th gen is NOT noticeably faster than the Core M powered machine it replaced."*

Similar comparisons are made for other aspects where customers draw parallels between the features of the reviewed product and those of their previous devices.

Variants Customers frequently compare different configurations or versions of the same laptop model, highlighting variations in storage capacity or RAM size. For example, a customer may discuss the availability of multiple storage options for a particular laptop model, such as a 128GB variant and a 64GB variant, or different processors:

> *"I did almost buy the amd version because the processor and gpu benchmark slightly higher but I didn't want to spend extra to get the ssd version of it."*
> *"Some of the 64gb versions appear to be the same as the 128gb laptops (check the specs), at a significantly cheaper price"*

Desired Product Some customers discuss their needs of laptop specifications rather the actual values of the product in question. Those desired, ideal values are different from the actual value of the product, leading to incorrect data in extracting aspect values. This phenomenon was observed across all aspects, for example:

> *"I got 8gb and wanted 16gb but wasn't going to pay the ridiculous cost difference"*
> *"need to order a screen replacement to have the laptop upgraded to full hd as I wanted."*
> *"I recently purchased the X1 after researching several other models (HP Spectre, Dell XPS)"*
> *"I needed a laptop that was budget friendly (under $500)"*

The phenomenon of cross-product references complicates the task of isolating and accurately assessing the value proposition of a specific laptop. For each utterance, a mechanism is needed that determines if the utterance talks about the current product, or is a cross-reference to another product (previous, variant, ideal).

4.2. Lesson 2: Varied Value Formatting

In our analysis, we noticed that customers write down information in their own style and do not necessarily make use of standard units and style conventions. For example, some customers occasionally employ informal, spoken language, such as abbreviations and synonyms:

> *"HP"*, *"h-p"* to refer to Hewlett-Packard
> *"GB"*, *"gb"*, *"gigs"*, *"gigabytes"* to refer to gigabytes

Additionally, customers employ symbols to represent product attributes or units:

> *"a duel channel kit of Corsair vengeance @ 2933"* signifying a speed of 2.933 MHz
> *"would recommend spending some more $$ for a better display"*
> *"3-*"* signifying the star rating
> *"DEFECTIVE ON #5"* signifying number of months

Formatting problems in user reviews have been observed in earlier works, showing how different styles of online reviews influence the evaluation process for the review writer [25] or leads to problems for automated methods [17, 18]. For humans, it serves the ease of writing or

ease of reading, but for extracting by algorithm, more elaborate coding schemes are needed, e.g., an additional dictionary-based method for standardization [26].

4.3. Lesson 3: Maximum Values instead of Real-Life Values

Another discrepancy we observed stems from sellers advertising the maximum capacities of their products, while reviewers focus on the experienced, real-world values. While the storage of a product is, at its full capacity, 512GB, some parts of it are reserved for the operating system, leaving less storage space at free disposal:

> *"at "256gb" you realistically only get 180 because windows takes about 70-80 gigs"*
> *"SSD memory is LESS than advertised in description (actual memory is 452 GB vs 512 GB)"*

Similarly, customers report a battery life that is less than the battery life stated by sellers:

> *"Battery life on 100% will give you 4-13 hours [...] NOT 15 hours as they claim"*

Contextual factors such as usage patterns play a role in the experienced and subsequently reported battery life. For the same laptop, customers with different usage patterns report different values (both examples were written for the same laptop):

> *"I wanted something that I could use to play Deep Rock Galactic, and this seemed to be the cheapest gaming laptop [...] the battery life only seemed to be 1-1.5 hours"*
> *"The battery life gives you about 5-6 hours of [...] internet browsing, MS office, emails"*

4.4. Lesson 4: Emphasis on Expandability

Customer reviews frequently discussed possible or performed upgrades of the products, such as upgrading the RAM or the storage capacity. These occurrences reflect customers' desires for enhanced performance and functionality and adapting the product to their individual needs. These mentions provide insights into evolving customer needs and preferences, but also lead to occurrences of values that do not overlap with the current product information. For example:

> *"Only thing I did was add another SSD (480GB Kingston 2.5) and more ram (Samsung 8gb)"*
> *"I recommend [...] upgrade their ram to dual channel 24 or 32gb ram."*

Both behaviors, talking about the possibilities for upgrades and the personal needs for upgrades, result in utterances where values are reported that the product does not (yet) have. If you consider the second example, the customer mentions two RAM values that are not the actual values of the laptop but upgrade options. Consequentially, automated extraction methods might pick up incorrect values.

4.5. Lesson 5: Approximate Value Mentions in Product Reviews

Another notable source of discrepancy is the tendency of customers to mention approximate values rather than exact figures in their product reviews. For instance, when discussing storage capacity, they often use rounded numbers such as *"250GB"* instead of specifying the exact capacity (in this case: 256GB). This trend extends across various product aspects, including processor speed, battery life, display resolution, and price. While some customers explicitly denote their approximation (e.g., *"For a brand new computer that cost almost $500 should work better."*, others use approximate values without signaling it in their text:

"only 250GB Storage isn't really enough to run and install most software (like games)"
"I received the product on Monday and it died on Friday. It is now a $300 doorstop."

In literature on vagueness in natural language, approximate values are a known strategy to reduce physical and mental effort [27]. If preciseness is not necessary for the main message, approximations can be used. For example, "$300" can be easier to type and to understand than "$286.99". For the reviewers, the precise storage capacity might not be important when assessing the laptop as a whole. Automated extraction methods need to take these strategies into account when extracting product information from natural language texts.

4.6. Additional Lesson: Price Values

One additional challenge that we observed was specific to the price aspect. Not only did the customers make use of approximate values (see Lesson 5), the price ground truth was fluctuating even during the course of the analysis. E-commerce platforms often employ dynamic pricing strategies [28], meaning that price values mentioned in review texts are likely outdated, even when mentioning rather precise values (e.g., *"Got it for $407"*, *"The Lenovo price on Amazon is usually $619, the Acer $749"*). Some customers also mention specific deals (e.g., *"I got this for $699 on cyber Monday"*). E-commerce platforms sometimes aggregate offers from various sellers or offer refurbished products at a lower price point, which might be another source of varying price values in reviews. Due to those factors, the values mentioned by customers in their reviews often diverge from the actual value of the product. Similarly, qualitative statements about prices from reviews (e.g., *"the price was very reasonable"*, *"the price is great"*, or *"Laptop is very good for the price"*) do not hold true over time. Practitioners and researchers who use reviews to extract information about products should carefully consider whether the information they target is dynamic and fluctuating over time – be it as dynamic as pricing information, or slowly changing over time like what customers consider to be "good" performance [29].

5. Conclusion

Customer-written review texts from e-commerce platforms (e.g., *Amazon*, *BestBuy* or *Alibaba*) are used as a data source for automatic tasks ranging from aspect extraction for summarization and market research to serving as signal impacting the product ranking. This research investigates the potential of reviews to serve as reliable source of information about products, focusing on the use case of laptop reviews. We annotated 1500 reviews and compared how well they reflected the true features of the laptops. Our analyses found that review texts contain inaccurate information about the reviewed product, refer to different product and their specifications, and misleading information about potential alterations and upgrades to the product. Using review texts for information extraction without accounting for those phenomena could therefore lead to inaccurate results. With this research, we contribute insights that can help e-commerce platforms and researchers to study online shopping habits in the future. Our research could also inform future automatic methods utilizing review data such as large language model deployment for e-commerce.

Acknowledgments

References

[1] M. Kovacs, S. Farias, F. Moura, A. Souza, Relations between consumer effort, risk reduction strategies, and satisfaction with the e-commerce buying process: the development of a conceptual framework, International journal of management 28 (2011) 316.

[2] T. Chen, P. Samaranayake, X. Cen, M. Qi, Y.-C. Lan, The impact of online reviews on consumers' purchasing decisions: Evidence from an eye-tracking study, Frontiers in Psychology 13 (2022) 865702.

[3] E. Constantinides, N. I. Holleschovsky, Impact of online product reviews on purchasing decisions, in: 12th International Conference on Web Information Systems and Technologies, WEBIST 2016, SCITEPRESS, 2016, pp. 271–278.

[4] Y.-C. Lien, R. Zhang, F. M. Harper, V. Murdock, C.-J. Lee, Leveraging customer reviews for e-commerce query generation, in: European Conference on Information Retrieval, Springer, 2022, pp. 190–198.

[5] X. Li, L. M. Hitt, Z. J. Zhang, Product reviews and competition in markets for repeat purchase products, Journal of Management Information Systems 27 (2011) 9–42.

[6] Precedence Research, E-commerce market (by model type: Business to business (b2b), business to consumer (b2c); by application: Home appliances, fashion products, groceries, books, others) - global industry analysis, size, share, growth, trends, regional outlook, and forecast 2023-2032, 2023. URL: https://www.precedenceresearch.com/e-commerce-market#:~:text=The%20global%20e%2Dcommerce%20market%20size%20exceeded%20USD%2014.14%20trillion,15%25%20between%202023%20and%202032.

[7] G. Askalidis, E. Malthouse, The value of online customer reviews, 2016. doi:10.1145/2959100.2959181.

[8] X. Zheng, S. Zhu, Z. Lin, Capturing the essence of word-of-mouth for social commerce: Assessing the quality of online e-commerce reviews by a semi-supervised approach, Decision Support Systems 56 (2013) 211–222.

[9] T. Melo, A. Silva, E. Moura, P. Calado, Opinionlink: Leveraging user opinions for product catalog enrichment, Information Processing and Management 56 (2019) 823–843. doi:10.1016/j.ipm.2019.01.004.

[10] Z.-J. Zha, J. Yu, J. Tang, M. Wang, T.-S. Chua, Product aspect ranking and its applications, IEEE transactions on knowledge and data engineering 26 (2013) 1211–1224.

[11] S. Chaabna, W. Hu, M. Lutf, Designing a ranking system for product search engine based on mining ugc, J. Manag. Inf. Syst. E-Commerce 2 (2015) 23–65.

[12] E. Najmi, K. Hashmi, Z. Malik, A. Rezgui, H. U. Khan, Capra: a comprehensive approach to product ranking using customer reviews, Computing 97 (2015) 843–867. URL: https://doi.org/10.1007/s00607-015-0439-8. doi:10.1007/s00607-015-0439-8.

[13] F. Sabbah, N. Fuhr, A transparent logical framework for aspect-oriented product ranking

based on user reviews, in: D. Hiemstra, M.-F. Moens, J. Mothe, R. Perego, M. Potthast, F. Sebastiani (Eds.), Advances in Information Retrieval, Springer International Publishing, Cham, 2021, pp. 558–571.

[14] J. Feuerbach, B. Loepp, C.-M. Barbu, J. Ziegler, Enhancing an interactive recommendation system with review-based information filtering, in: P. Brusilovsky, M. de Gemmis, A. Felfernig, P. Lops, J. O'Donovan, N. Tintarev, M. C. Willemsen (Eds.), Interfaces and Human Decision Making for Recommender Systems: Proceedings of the 4th Joint Workshop on Interfaces and Human Decision Making for Recommender Systems co-located with ACM Conference on Recommender Systems (RecSys 2017), volume 1884 of *CEUR workshop proceedings*, RWTH Aachen, Aachen, 2017, pp. 2–9. URL: http://ceur-ws.org/Vol-1884/paper1.pdf, oA platinum.

[15] Y. Wang, L. Luo, H. Liu, Bridging the semantic gap between customer needs and design specifications using user-generated content, IEEE Transactions on Engineering Management 69 (2022) 1622–1634. doi:10.1109/TEM.2020.3021698.

[16] S. Novgorodov, I. Guy, G. Elad, K. Radinsky, Generating product descriptions from user reviews, in: The World Wide Web Conference, WWW '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 1354–1364. URL: https://doi.org/10.1145/3308558.3313532. doi:10.1145/3308558.3313532.

[17] B. Seerat, F. Azam, Opinion mining: Issues and challenges(a survey), International Journal of Computer Applications 49 (????) 42–51.

[18] A. Valdivia, E. Hrabova, I. Chaturvedi, M. V. Luzón, L. Troiano, E. Cambria, F. Herrera, Inconsistencies on tripadvisor reviews: A unified index between users and sentiment analysis methods, Neurocomputing 353 (2019) 3–16.

[19] M. Hu, B. Liu, Mining and summarizing customer reviews, in: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04, Association for Computing Machinery, New York, NY, USA, 2004, p. 168–177. URL: https://doi.org/10.1145/1014052.1014073. doi:10.1145/1014052.1014073.

[20] R. M. D'Addio, M. A. Domingues, M. G. Manzato, Exploiting feature extraction techniques on users' reviews for movies recommendation, Journal of the Brazilian Computer Society 23 (2017) 1–16.

[21] A. Yates, J. Joseph, A.-M. Popescu, A. D. Cohn, N. Sillick, Shopsmart: product recommendations through technical specifications and user reviews, in: Proceedings of the 17th ACM conference on Information and knowledge management, 2008, pp. 1501–1502.

[22] H. Li, P. Yuan, S. Xu, Y. Wu, X. He, B. Zhou, Aspect-aware multimodal summarization for chinese e-commerce products, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 8188–8195.

[23] S. Jain, P. Hegade, E-commerce product recommendation based on product specification and similarity, in: 2021 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT), IEEE, 2021, pp. 620–625.

[24] V. Braun, V. Clarke, Reflecting on reflexive thematic analysis, Qualitative Research in Sport, Exercise and Health 11 (2019) 589–597. URL: https://doi.org/10.1080/2159676X.2019.1628806. doi:10.1080/2159676X.2019.1628806. arXiv:https://doi.org/10.1080/2159676X.2019.1628806.

[25] J. Grigsby, C. Zamudio, R. Jewell, Not so bad after all: How format influences review writers'

post-review evaluations, Journal of Consumer Behaviour 21 (2022). doi:10.1002/cb.2079.

[26] H. V. Cook, L. J. Jensen, A guide to dictionary-based text mining, Bioinformatics and drug discovery (2019) 73–89.

[27] S. Solt, Vagueness and imprecision: Empirical foundations, Annual Review of Linguistics 1 (2015) 107–127. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-linguist-030514-125150. doi:https://doi.org/10.1146/annurev-linguist-030514-125150.

[28] J. Nicas, Now prices can change from minute to minute, Wall Street Journal (2015).

[29] J. Liu, Y. Zhang, X. Wang, Y. Deng, X. Wu, Dynamic pricing on e-commerce platform with deep reinforcement learning: A field experiment, arXiv preprint arXiv:1912.02572 (2019).