A visual tool for reducing returns in e-commerce platforms

Riccardo Cantini
DIMES, University of Calabria, Italy
rcantini@dimes.unical.it

Fabrizio Marozzo
DIMES, University of Calabria, Italy
fmarozzo@dimes.unical.it

Alessio Orsino
DIMES, University of Calabria, Italy
aorsino@dimes.unical.it

Martina Passarelli
DIMES, University of Calabria, Italy
mpassarelli@dimes.unical.it

Paolo Trunfio
DIMES, University of Calabria, Italy
trunfio@dimes.unical.it

***Abstract*—Nowadays, the number of people who prefer to make online purchases on e-commerce platforms is constantly increasing. Online shopping turns out to be fast and cheap but it also involves some risks. In fact, it sometimes happens that a product, once received, is smaller or bigger than expected, or does not comply with what is described on the purchase site. Typically this happens because the information provided by retailers is not detailed and, moreover, the images usually show the product on a white background without any reference nearby to understand its true dimensions. To cope with this problem, we developed *CompareDimensions*, a software tool that helps customers during their online purchases, by allowing them to understand the correct dimensions of the product they are about to buy. Through this tool, a customer can see the images of the product and compare it with a series of reference images known to all, such as a man, a chair, a door or a child. The experimental evaluation, conducted on Amazon, shows that the use of the proposed tool helped people to figure out the real dimensions of products, with respect to the case in which only cut-out images and textual descriptions are provided. In particular, according to our experiments, people who used the tool were able to understand the real dimensions of a product with 17.5% greater accuracy than those who did not use the tool. Focusing on critical products, average accuracy increased by 24%, with an improvement of 27% and 11% for products that proved to be smaller or larger than expected. Therefore, the adoption of CompareDimensions could help customers to understand the real dimensions of the products they want to buy, and consequently reduce the likelihood of a return.**

***Index Terms*—E-commerce, Online selling, Returns, Waste, Amazon, Product dimensions.**

I. Introduction

E-commerce, also known as electronic commerce, refers to transactions conducted via the Internet. The term is often used to refer to the sale of physical products online, but it can also describe any kind of commercial transactions facilitated through the Internet. Leading e-commerce platforms that provide Business to Consumer (B2C) services, such as Amazon, eBay, Walmart, Taobao, and Alibaba, allow customers to buy and sell products globally, at any time and in any place, usually with advantageous prices due to reduced distribution costs and various promotions. Other important advantages compared to physical stores are the wide choice of products, the possibility to read reviews of other customers, and above all clear return policies which in most cases include free returns.

Indeed, returns management is crucial since it is one of the main factors that determines the choice of a particular online shopping website by a customer. Easy returns and fast refunds are helpful in winning customer loyalty in the e-commerce business. However, along with the increase in CO_2 emissions and underutilized transport systems for single deliveries, returned products are one of the main problems of online shopping [1]. Globally, about 36% of online shoppers return some or all their online orders [2], and less than half of returned goods are re-sold at full price. This shift in consumer behavior has forced retailers to produce more than the real demand, with the consequence of reduced margins because of costs, customer distrust, and severe environmental impacts. Within the EU alone, around three billion purchases are returned to the retailers per year, with huge economic and environmental costs due to customer service, financial reconciliations, warehousing, restocking, repackaging, remarketing, and shipping services.

The main reasons behind returns are summarized in Figure 1, which shows data collected in a report provided by *UPS* [2], the world's largest package delivery company. The report is the result of a survey administered to 18,000 global shoppers between December 2018 and January 2019.

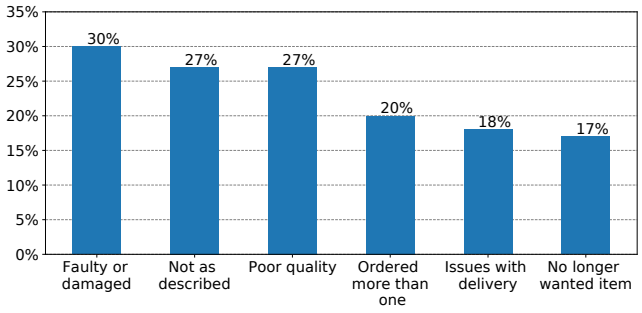

Fig. 1: The main reasons behind product returns according to UPS company.

Leaving out the faulty or damaged products, which necessarily require a return to be replaced, the second most common reason for returns is that the product is different from the textual description and/or images provided by retailers. In fact, online shoppers can only rely on product descriptions, images, and reviews of other customers. However, different marketing techniques can be used by retailers to highlight product properties making it more attractive. This can lead to customers returning products deemed to be different from those expected. This reason includes those scenarios in which the dimensions of the product can be misleading, as not all people can figure them out from the technical description.

For all these reasons, we designed and developed *CompareDimensions*, a software tool that aims aiding the customers in their online shopping, allowing them to understand the real dimensions of the products they are about to purchase. Using this tool, returns due to unexpected or wrong dimensions could be significantly reduced, with great advantages from an economic and environmental point of view. Customers can compare the images of the product with different references, selected among those that are considered more familiar to the human brain (e.g., a man, a child, a chair, a door). Product images can be moved, brought closer to reference images, and measured in a 2D environment using a Cartesian plane whose dimensions can be shown in different units of measurement. This can help reduce the reasons for making a return, which in turn results in less economic losses for retailers and a reduced impact for the environment.

The effectiveness of the proposed tool was assessed using Amazon as a testbed. In addition to being the most widely used e-commerce site, an UPS report [2] highlights that 56% of the online shoppers start their search directly on Amazon rather than using search engines (Google, Bing, Baidu, etc.), even when they do not buy there. The experimental evaluation, conducted through a survey, showed that the use of CompareDimensions helped customers to figure out the real dimensions of a product. In particular, the number of correct answers about the true dimensions of a product increased by 17.5% with the aid of the proposed tool. Moreover, by focusing on critical products, i.e. characterized by negative customer reviews related to their misleading dimensions, we observed an improvement of 24% in average accuracy. Specifically, the tool has brought the greatest advantages in the case of products whose size was smaller than expected, with 27% of more correct answers. Finally, we evaluated the benefits arising by the use of the tool with respect to product categories, such as heating and cooling ($+57\%$), home decor ($+35\%$), vacuum cleaner ($+35\%$), furniture ($+33\%$), and waste baskets ($+14\%$). These results show that our tool is effective in helping customers to correctly intepret the actual dimensions of a product, thereby reducing returns made for that cause.

The remainder of the paper is organized as follows. Section II discusses related work. Section III describes the proposed tool. Section IV presents the experimental evaluation and the main results, and Section V concludes the paper.

II. Related work

With the growth of the e-commerce business, more and more researchers focused on the return behavior of consumers and the problem of returns, providing beneficial return strategies to retailers [3]. In particular, economic studies are mainly oriented towards the mathematical formulation of optimal return policies [4], [5], [6] and the identification of measures for reducing the return rates [7], [8].

From the customer viewpoint, the return issue due to unexpected dimensions has been addressed by some tools, websites, and plug-ins. Some of the most important are discussed in the following, highlighting strengths and weaknesses.

PiliApp [9] is a website that provides a huge set of product images including smartphones, tablets, and more. Objects are displayed in full-size and resized according to the screen characteristics. Even if the database is constantly updated, the site does not handle external products. Moreover, objects are not catalogued and it is impossible to search by category.

Pective [10] is a website very similar to PiliApp, but with some improvements. In particular, users can upload new images in full scale and evaluate the correctness of the image resizing with respect to the screen.

Compare Sizes [11] is a website that allows comparing multiple objects by manually specifying their dimensions. Objects are displayed side by side, offering a proportionate view of their dimensions. It also provides everyday items and a wide catalog of commercial products, but objects are represented only as boxes, which may give a distorted perception of relative dimensions during the comparison.

Phone Sized [12] is a website for comparing the relative dimensions of mobile phone and tablet devices. Compared to the other analyzed solutions, this appears to be very sectorial.

Finally, *Next to a person* [13] is another plug-in, also available as a Google Chrome extension, that only works for the Amazon site. It scans an Amazon page extracting the dimensions, takes the main image and resizes it compared to an average sized human, which is $1.7m$ for a man and $1.6m$ for a woman. However, in most cases the tool is unable to return any valuable result, due to the lack of variety in reference images for making the comparisons.

All the aforementioned tools are aimed at supporting the customer in interpreting correctly the dimensions of a product they are going to buy, even if they present a series of limitations as described above. The tool presented in this paper, CompareDimensions, is able to overcome these issues, and provides an effective and easy way to give to the customer a clear idea about product dimensions. In particular, differently from [9], [10] and [12], the tool shows a side by side comparison with different reference products within a reference system represented by two Cartesian axes, which is useful to represent real dimensions. Compared to [13], instead, our tool exploits all the images provided by the retailer and relies on a varied set of reference products, leading to more meaningful comparisons. Finally, unlike [11], CompareDimensions shows products in their original form, without replacing them with boxes which can distort the perception of relative dimensions.

III. CompareDimensions

CompareDimensions is aimed at supporting the purchase process of an e-commerce customer, so as to avoid a return due to unexpected or poorly described dimensions. Through this tool, customers can compare a product with a range of references chosen from a database in order to have a clear and precise idea of the dimensions of what they are purchasing. The tool has been successfully tested on Amazon, but its usage can be extended to all e-commerce platforms. With CompareDimensions, a user-friendly visual interface shows the product from the e-commerce page next to one of the references in Table I. The user can change the reference and the product image to make the necessary comparisons.

Name	Height (cm)	Length (cm)
man	175	54
woman	165	45
teenager	150	40
child	100	32
newborn	50	50
school chair	95	50
kitchen chair	82	40
door	210	80

Tab. I: Main references for product comparison.

An e-commerce page can be parsed as a HTML document, which stores product dimensions and image information. Generally, the dimensions of a product are displayed in the first box reserved to its description; otherwise, they are displayed in the technical details area. In addition to the numeric dimensions, retailers usually provide a large number of images of the product, from different positions and angles, to give a complete view of it. Generally, among all the provided images, the first one shows the product from the front, as in Figure 2. Nevertheless, when an object is shown without a known reference nearby, it can be difficult for the customer to understand its true dimensions.

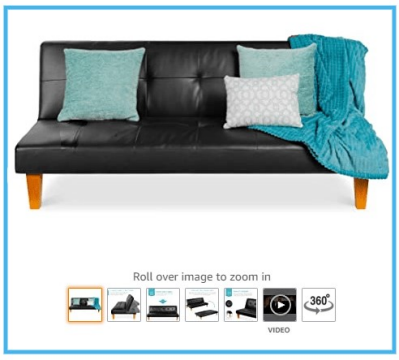

Fig. 2: Product images by the retailer (Amazon.com).

CompareDimensions acts in three main steps:

1) It extracts and validates the dimensions of the product and the pattern to be used to interpret them.
2) It collects and processes the provided images.
3) It compares the product images with well-known references in a graphic panel, showing them side by side. The images are resized to provide a proportionate view and the user is given the possibility to add Cartesian axes with measures in several formats (e.g., centimetres, inches, etc.). Additionally, it is possible to manually modify product dimensions in order to deal with any error in the textual description provided of the retailer.

ALGORITHM 1: Find the dimension pattern

Input : An HTML document
Output : An array with dimensions sorted according to $L \times W \times H$

```
Initializing regex_number_and_units, regex_units, aid_array
  and result variables;
first_dimension_found=first_research();
for element in document do
    if the text of the element matches with regex_units then
        add the matches text in aid_array;
Remove duplicates from the aid_array;
for element in aid_array do
    if the element matches with regex_number_and_units
      and includes first_dimension_found then
        result=element;
        break for;
if result is an empty string then
    return
      process_dimension_default(first_dimension_found);
else
    return process_dimension(result);
```

Algorithm 1 implements the first step, allowing the extraction of the information about product dimensions on a generic web page. Usually, the most common pattern adopted to describe the dimensions is $n_1 \times n_2 \times n_3$, which can be preceded by a literal pattern that suggests the interpretation of those measures. If no literal pattern exists in this way, the dimensions will be interpreted according to the default pattern, i.e. $Length \times Width \times Height$ ($L \times W \times H$). Otherwise, they will be interpreted as stated by the found pattern. The proposed procedure performs an analysis of the textual content of the page and exploits a set of regular expressions to find and validate the dimensions of the product. Finally, they are saved and returned in an array sorted according to the $L \times W \times H$ format. The procedure has been successfully tested on structured e-commerce sites such as Amazon, eBay, and Alibaba. In particular, for each site we selected a set of 100 products, measuring the percentage of products for which dimensions were successfully extracted (i.e., the hit rate). Achieved results are summarized in Table II.

E-commerce	Total products	Hit rate
Amazon	100	97%
eBay	100	95%
Alibaba	100	94%

Tab. II: Hit rate of Algorithm 1 for different e-commerce sites.

There exist two main scenarios in which the extracted dimensions may be wrong, which can lead to poorly representative comparisons:

- The pattern for interpreting the dimensions is completely absent or it does not follow the default one. In this case, the error is due to a lack in the description.
- The package dimensions are specified in place of product ones, and the difference may be not negligible.

The second step is to collect all the images provided by the retailer for a given product. They are available in different resolutions, but we have kept only the high resolution ones. Product images present on popular e-commerce usually have a white background and often an outer frame, which can lead to an incorrect resizing. Therefore, the tool exploits a procedure to remove the background to get only the cropped representation of the product. In particular, it iterates the pixel matrix of the image identifying the background regions to be removed in order to obtain the minimum sized frame that contains the product. Finally, images are resized according to the found dimensions and aggregated in a graphic panel to be compared with different references loaded from Table I. The entire process is detailed in Algorithm 2.

ALGORITHM 2: Compare dimensions

Input : An HTML document, the dimensions of the product (LxWxH)
Output : Interactive panel with images compared

```
Initializing array of product_images;
for image in document do
    if the resolution of image is LARGE then
        image = remove_background(image);
        product_images.add(image);
for image in product_images do
    image = resize_image(image, dimensions);
reference_images = load_reference_images();
panel.show(reference_images, product_images);
return panel;
```

The final result is shown in Figure 3. In particular, it highlights the use of the tool to analyze the dimensions of the sofa shown on Figure 2. In the example, the reference on the left is a man of standard dimensions but the customer can navigate between different references (e.g., a child, a chair, a door, and more) and compare them with the images provided by the retailer. The comparison is inserted in Cartesian axes and a conversion button between measurements is provided.

IV. Experimental evaluation

This section is aimed at presenting and discussing the main benefits related to the use of CompareDimensions in the purchasing process. The effectiveness of the proposed tool was assessed through an extensive evaluation using Amazon as a testbed. In particular we realized an implementation of the proposed tool by developing a Google Chrome extension which works on Amazon websites, available at https://github.com/SCAlabUnical/CompareDimensions. Experimental evaluation was carried out in three main steps, as described in the following.

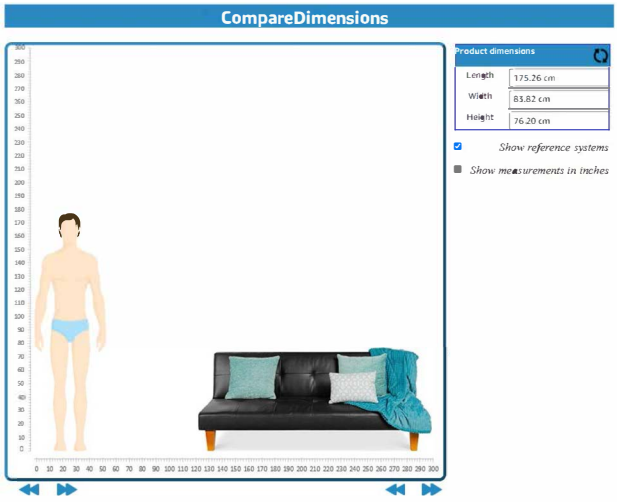

Fig. 3: Example of using CompareDimensions to analyze the dimensions of a sofa.

- A set of 100 products was selected from the Amazon website, in order to administer a survey in which a group of 100 users were asked about the dimensions of some products with and without the aid of the tool.
- The advantages provided by the use of CompareDimensions and their statistical significance were investigated, together with the behavior of the tool in relation to the different product categories.
- The products used in the survey were classified in order to analyze the main reasons behind a return (e.g., a product is too small or too big with respect to the expected dimensions) along with the benefits brought by the tool in the presence of each of these issues.

A. Survey administration and results

As a first step, a set of 100 products, sold in a period from 2017 to 2021, was randomly selected from Amazon.com. Afterwards, a survey was prepared, composed of 20 questions where the users were asked to interpret the dimensions of a product, by a comparison with well-known references from everyday life. The references used in the survey are detailed in Table III.

Name	Height (cm)	Length (cm)
wine bottle	30	8
mug	12	16
single bed	190	80
A4 sheet paper	30	21
blackboard	100	200
door	210	80
Mona Lisa	77	53
kitchen chair	82	40
1 KG pack of pasta	31	21

Tab. III: References used in the administration of the survey.

Each reference to be associated with a product has been chosen as the one with the most similar dimensions in order

to make the comparisons more complex for the testers. The survey was submitted to a sample of 100 users, evaluating their ability to correctly figure out the real dimensions of the different items with and without the aid of the tool. Questions were asked in the form "Do you think product x is wider/higher than reference y?". Figure 4 shows a first result from the survey. In particular, out of 2000 total questions, we registered an increment of the number of correct answers of 17.5% using CompareDimensions (76.5% versus 59%).

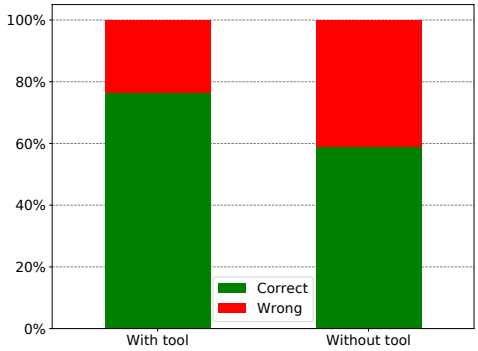

Fig. 4: Average accuracy of dimension assessment based on the answers of the surveyed users.

However, although we observed a higher average accuracy achieved by using the proposed tool, this result is not enough to state that this improvement is statistically relevant. For this reason, we used a statistical test based on t-Student distribution (i.e., t-test), for evaluating whether the differences registered with and without the use of CompareDimensions is either due to chance, or to an actual advantage brought by the tool. In order to perform the aforementioned test, we considered the accuracy of each user's assessments in the two scenarios (i.e., without and with the aid of the tool), evaluating the statistical significance in the increase of average accuracy with a level of significance $\alpha = 0.05$. T-test is mostly used when data follows a normal distribution and may have unknown variances. Therefore, we exploited the Shapiro-Wilk test for checking if user accuracy, both without and with the use of CompareDimensions, is normally distributed. In particular, we obtained that both distributions are normally distributed and mesokurtic. Subsequently, we performed a one tailed t-test on paired data to compare the average user accuracy in the two scenarios. Specifically, we found that the increase in average accuracy related to the use of CompareDimensions is significant for $p = 1.39 \times 10^{-5} < \alpha$, which supports the effectiveness of the proposed tool.

We also analyzed how the category of the different products involved in the survey influenced the average accuracy of the users. The considered products can be grouped in 14 different categories: *furniture* (chairs, tables, armchairs, sofas), *kitchen devices* (refrigerators, toasters, ovens, kitchen robots), *home decor* (mirrors, lights, lamps), *bookcase*, *plant accessories* (pots, pot holders, plant stands), *crib*, *organizer* (toy box, clothes box, tool box), *heating and cooling* (stoves, hot and cold air devices), *vacuum cleaner*, *TV and PC accessories* (PC screens, televisions, printers), *closet system* (cabinets, hangers), *pet accessories* (aquariums, animal cages, cat toys), and *toys*. Figure 5 shows the top-5 categories in which the tool has helped users the most to have a correct perception of product dimensions. Specifically, we observed the following improvements on average accuracy: 57% for heating and cooling, 35% for home decor, 35% for vacuum cleaner, 33% for furniture, and 14% for waste basket.

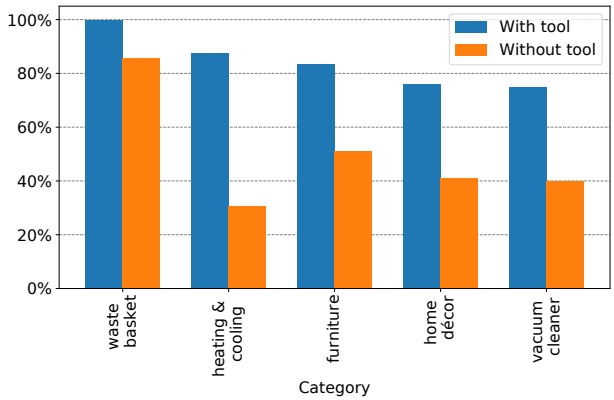

Fig. 5: Average accuracy per product category.

B. Critical product analysis

In this section we provide an analysis of the advantages brought by the proposed tool in relation to the main reasons behind a return. In particular, we focused on the two main scenarios related to a return caused by a wrong perception of product dimensions, i.e the product is *smaller* or *bigger* than expected. We exploited a keyword-based classification which relies on customer reviews for labeling products as follows:

- *Smaller* or *bigger*, if the majority of customers who have purchased a product complain that its dimensions do not conform to the description provided by the retailer.
- *Unknown*, if no complaint reviews have been extracted for that product. This means that there is no evidence that its dimensions are misleading.

The classification process exploits a set of keywords $\mathcal{K}$ generated from a series of adjectives that qualify the dimensions of a product, i.e. *small*, *short*, *tight*, *big*, *wide*, and *large*, following different patterns. Specifically, we combined these adjectives with intensifiers like *too*, *very*, *extra*, and *super*, searching also for expressions like "*smaller/shorter/bigger/*...than expected". Moreover, in order to handle the negation, we used other patterns such as "not small/short/big/...enough". The keywords in $\mathcal{K}$ are used to extract negative comments for a given product p, which is then labeled as *smaller* or *bigger* if most of the complaints fall into that class. If no comments are extracted for p, it is classified as *unknown*. Table IV shows some examples of negative comments from Amazon, along with the class label assigned to the related product.

Using this classification process, we labeled 64 products as *smaller*, 14 as *bigger*, and 22 as *unknown*. We collected in

Comment	Class label
Returned. **Too big**	bigger
Item is far **too large**.	bigger
The base is **too wide** to fit in a smallish space.	bigger
Too small, sent it back.	smaller
The lampshade is **super short**! Not as pictured.	smaller
Was **too small** returned item.	smaller

Tab. IV: An example of negative comments along with the label of the related product (Amazon.com)

a dataset all the information about each product, including its category, the hyperlink to Amazon, the assigned label (*smaller*, *bigger* or *unknown*), the most critical comment, its description and dimensions. This dataset is openly available at https://github.com/SCAlabUnical/CompareDimensions/.

Once all products were labeled as described above, we focused our analysis only on critical products, i.e. those with a declared issue related to a misinterpretation of their dimensions. Thus, we investigated the benefits brought by the use of CompareDimensions when a given product is *smaller* or *bigger* than expected. The increment in average accuracy for the two scenarios are shown in Figure 6.

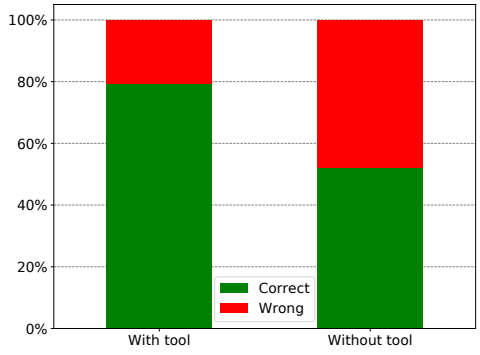

(a) Results for class *smaller*.

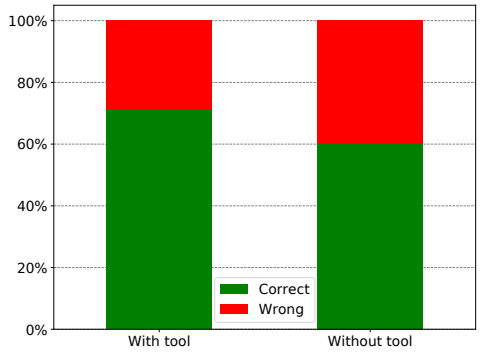

(b) Results for class *bigger*.

Fig. 6: Average accuracy of dimension assessment by the surveyed users focusing on critical products.

Overall, for products categorized as critical, we obtained an improvement of 24%. Specifically, we observed a greater improvement for products of class *smaller*, compared to those labeled as *bigger*, with an increase in average accuracy of 27% and 11% respectively. This result shows the tendency of customers to overestimate product dimensions. In fact, when only traditional cut-out images are provided, just showing the product on a white background, customers often misinterpret its real dimensions. In particular, they are more likely to consider products to be greater than they actually are, since images are provided without a scale and smaller products are often zoomed to make the details more visible and clear.

V. Conclusions

The use of e-commerce platforms has grown a lot in the last few years, thanks to many factors including a greater convenience and variety of products compared to physical stores and advantageous return policies. However, users often misinterpret the real dimensions of the product they are going to buy, which can cause a return. This scenario leads to increased costs for retailers and a waste of time for customers, along with a series of negative effects on the environment. In order to overcome these issues, we proposed *CompareDimensions*, a tool designed to help customers during their online purchases, allowing them to correctly estimate the dimensions of a product by comparing it with a wide variety of everyday life reference images, such as a man, a chair, or a door. The tool was tested using Amazon as a testbed, by conducting a survey on 100 users. Experimental results showed its effectiveness in helping customers to get a clear idea of the actual dimensions of a product. We found that the number of correct answers about the true dimensions of a product increased by 24% for critical products. Specifically, the impact of the tool was higher for products whose actual size was smaller than expected, with 27% more correct answers, due to customers' tendency to overestimate real product dimensions.

References

[1] Alexander K Seewald, Thomas Wernbacher, Alex Pfeiffer, Natalie Denk, Mario Platzer, Martin Berger, and Thomas Winter. Towards minimizing e-commerce returns for clothing. In *ICAART (2)*, pages 801–808, 2019.

[2] UPS. Ups pulse of the online shopper. https://solutions.ups.com/rs/935-KKE-240/images/UPS-Pulse-of-the-Online-Shopper-Report.pdf.

[3] Huaman Han. Review and prospect on return problems of e-commerce platform. *Open Journal of Business and Management*, 2019.

[4] Scott Davis, Eitan Gerstner, and Michael Hagerty. Money back guarantees in retailing: matching products to consumer tastes. *Journal of Retailing*, 71(1):7–22, 1995.

[5] Samar K. Mukhopadhyay and Robert Setoputro. Optimal return policy and modular design for build-to-order products. *Journal of Operations Management*, 23(5):496–506, 2005.

[6] Baris Yalabik, Nicholas C. Petruzzi, and Dilip Chhajed. An integrated product returns model with logistics and marketing coordination. *European Journal of Operational Research*, 161(1):162–182, 2005.

[7] Ximin Huang, Sin-Man Choi, Wai-Ki Ching, Tak-Kuen Siu, and Min Huang. On supply chain coordination for false failure returns: A quantity discount contract approach. *International Journal of Production Economics*, 133(2):634–644, 2011.

[8] Mark Ferguson, V. Daniel R. Guide, and Gilvan C. Souza. Supply chain coordination for false failure returns. *Manufacturing & Service Operations Management*, 8(4):376–393, 2006.

[9] PiliApp. https://www.piliapp.com/actual-size/. Accessed 30 July 2021.

[10] Pective. http://pective.com/. Accessed 30 July 2021.

[11] CompareSizes. https://comparesizes.com/. Accessed 30 July 2021.

[12] Phone Sized. https://phonesized.com/. Accessed 30 July 2021.

[13] Next to a person. https://nexttoaperson.com/. Accessed 30 July 2021.