

# Revamp: Enhancing Accessible Information Seeking Experience of Online Shopping for Blind or Low Vision Users

Ruolin Wang  
UCLA HCI Research

Zixuan Chen  
UCLA HCI Research

Mingrui “Ray” Zhang  
The Information School,  
University of Washington

Zhaoheng Li  
Department of Computer Science and  
Technology, Tsinghua University

Zhixiu Liu  
Computer Science Department,  
Stanford University

Zihan Dang  
UCLA HCI Research

Chun Yu  
Department of Computer Science and  
Technology, Tsinghua University

Xiang “Anthony” Chen  
UCLA HCI Research



*“shaped like a Cola Bottle,  
the logo is sleek,  
It is a flat pink with a few speckles  
of white mixed in the blue,  
sized for the disposable plastic  
water bottles”*

**A** A brief appearance description  
read out as alt-text

**Simple Modern Stainless  
Steel Vacuum Insulated  
Double-Walled Wave  
Bottle, 17oz - Ombre:  
Sweet Taffy**

- Vacuum Insulated: The Simple Modern Wave water bottle is double-walled and vacuum insulated, which keeps your favorite beverage hot or cold for hours
- Premium 18/8 Stainless Steel: Bottle is made from high-quality, food-grade 18/8 stainless steel and is coated in a durable finish
- Premium Copper Coated Insulation Layer: The exterior of the inner wall, in the vacuum sealed area, is coated in a thin layer of copper for added insulation
- Take it With You: Narrow mouth reduces spills and sloshing while you drink; 9oz

**B** what is the shape?

**Summary:**

- C** I got 23 related reviews, 22 are positive, the top related snippets are:
- convenient shape fits in the side pockets of my bag or book bag so I need never thirst again
  - shaped like a Cola Bottle and has a screw top lid with a good size opening for drinking
  - The shape is a bit awkward, in my opinion, but it is a great alternative to any high end water bottle
- 1 are negative, the top related snippets are:
- knock off 1 star because the shape may limit its use for other beverages

**D** Positive List

Negative List

**[Title]** Great Product!  
**[Content]** convenient shape fits in the side pockets of my bag or book bag so I need never thirst again

**[Title]** Awesome Cola Shape  
**[Content]** shaped like a Cola Bottle and has a

**[Title]** Great bottle, but shape might limit its use for some.  
**[Content]** knock off 1 star because the shape may limit its use for other beverages

**Figure 1: Revamp simplifies and reconstructs the original Amazon web page for Blind or Low Vision users’ information seeking task to understand product details, especially the appearance. Using rule-based heuristics, Revamp extracts descriptive information from customer reviews to generate image descriptions (a), responses to users’ queries (b) with a sentiment summary of all the relevant reviews (c) and original reviews sorted into a positive and a negative lists (d).**

## ABSTRACT

Online shopping has become a valuable modern convenience, but blind or low vision (BLV) users still face significant challenges using it, because of: 1) inadequate image descriptions and 2) the inability to filter large amounts of information using screen readers. To address those challenges, we propose Revamp, a system that leverages customer reviews for interactive information retrieval. Revamp is a browser integration that supports review-based question-answering interactions on a reconstructed product page. From our

interview, we identified four main aspects (color, logo, shape, and size) that are vital for BLV users to understand the visual appearance of a product. Based on the findings, we formulated syntactic rules to extract review snippets, which were used to generate image descriptions and responses to users' queries. Evaluations with eight BLV users showed that Revamp 1) provided useful descriptive information for understanding product appearance and 2) helped the participants locate key information efficiently.

## CCS CONCEPTS

• **Human-centered computing** → **Accessibility systems and tools**; **Natural language interfaces**.

## KEYWORDS

Online shopping, Information Retrieval, Accessibility, Blind or Low Vision Users, Reviews, Image Description, Question-answering

### ACM Reference Format:

Ruolin Wang, Zixuan Chen, Mingrui “Ray” Zhang, Zhaoheng Li, Zhixiu Liu, Zihan Dang, Chun Yu, and Xiang “Anthony” Chen. 2021. Revamp: Enhancing Accessible Information Seeking Experience of Online Shopping for Blind or Low Vision Users. In *CHI Conference on Human Factors in Computing Systems (CHI '21), May 8–13, 2021, Yokohama, Japan*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3411764.3445547>

## 1 INTRODUCTION

Online shopping has gained increasing popularity among Blind or Low Vision (BLV) people who have limited mobility to travel to physical stores. According to a research survey conducted in the UK, over 90% of people with disabilities shopped online at least once a month<sup>1</sup>. The recent COVID-19 pandemic further accelerated the adoption of online shopping among BLV users<sup>2</sup>. Thus making online shopping experience accessible has become an imperative requisite for ensuring the quality of life for BLV users.

However, prior studies [30, 31] show that BLV people still face significant barriers on online shopping websites due to inadequate image descriptions and screen readers' inability to filter out a large amount of information on a product page. Our formative study with 20 BLV people showed that automatic tools (e.g., Seeing AI, filters of screen reader) were only used by a few experienced users, and the information provided was too generic to inform a purchase decision, especially for certain categories where the appearance matters, e.g., fashion products. The most efficient way for BLV users is still seeking help from a sighted person, such as a family member or a crowdsourced helper, which is not always available.

To address the above challenges, we developed *Revamp*, an interactive information retrieval system that supports review-based question-answering (QA) on a reconstructed product page. It is implemented as a browser extension for Amazon.com as shown in Figure 1. Utilizing reviews as a knowledge source, Revamp extracts informative descriptions to help BLV users understand the product appearance. To understand what kinds of visual questions BLV users care about, we collected questions from ten BLV participants on their frequently shopping categories on Amazon (Clothing,

Shoes & Jewelry, Home & Kitchen, and Electronics), identified four main aspects (color, logo, shape, and size) in understanding the appearance. Based on the results, we formulated syntactic rules to extract the review snippets to generate image descriptions and responses to users' queries.

We tested the performance of these rules on the Amazon best-seller product lists covering three main shopping categories. Results showed that these rules (i) covered 85% informative reviews voted by BLV participants and (ii) provided insights on information conveyed by images evaluated by two sighted people. We evaluated the usability of Revamp with eight BLV people on six representative products in the three aforementioned categories. Participants reported that Revamp reduces their efforts of information seeking, improves their utilization of reviews, and extracts reviews that are informative to help them understand product appearance. They considered Revamp to be helpful for shopping online independently when no sighted helpers were available.

Our **contributions** are:

- Understanding the challenges faced by BLV people in information seeking when shopping online and design implications to meet their unique needs;
- Identifying specific questions important to BLV users when shopping online and deriving syntactic rules that retrieve informative reviews to provide image description or to answer visual questions;
- Revamp, an interactive information retrieval system that integrates with Amazon to support product browsing and review-based question-answering;
- Evaluation study with eight BLV users that validated the feasibility and usefulness of enhancing accessible online shopping experience via Revamp.

## 2 RELATED WORK

This research builds upon prior work from different sub-disciplines across accessibility, information retrieval (IR), computer vision (CV) and natural language processing (NLP). We first review previous research investigating online shopping experience for BLV users, as well as the efforts to improve the accessibility of such experience; we then provide an overview of information retrieval methods that utilize online reviews to generate useful information, which serves as the foundation of the implementation of Revamp.

### 2.1 Online Shopping Accessibility

Blind or low vision (BLV) people often take great effort to find and learn what products are visually appropriate for them [16, 29]. Without real-time communication with store clerks and touching the products directly, information related to online products is mainly conveyed through images and text provided by sellers. Hence, inadequate descriptions of images and unparsed text information significantly reduce BLV consumers' engagement with online shopping websites [30, 31]. Currently, BLV people could leverage human-powered assistive tools to answer visual questions about daily objects and products in stores [1–3, 6, 22, 27, 39]. Specially in shopping clothes, human helpers can provide remote assistant on describing articles of clothing (e.g., color) [1, 3] or offering

<sup>1</sup><http://www.clickawaypound.com/index.html>

<sup>2</sup><https://www.shropshirestar.com/news/health/coronavirus-covid19/2020/04/07/online-shopping-priority-plea-for-the-blind/>

subjective fashion advice [6, 22]. Relatedly, automatic photo caption generation and visual question answering in computer vision and AI research is often designed for general scenarios [8, 13, 26]. There exist automatic solutions that promote accessible image understanding on the web, e.g., providing image alt-text for photos in Facebook based on object recognition [38], using OCR for photos on Twitter and other websites alike [9, 18], and using reverse image search to find existing captions [10].

However, such solutions cannot fulfill BLV users' special need in better understanding visual details of online products. Most images on eCommerce websites, if captioned by emerging automatic caption tools such as SeeingAI<sup>1</sup>, will only have a generic description of the product, e.g., *"probably a close up of a red chair"*, without any specific insights on appearance details, e.g., how is the size, and how does this kind of red make people feel. BrowseWithMe [31] has made an important step forward in describing clothes outfit by identifying the image regions, but only respond with basic color names (e.g., *"Brick Top, Navy Pants"*) based on color detection of image. Our contribution is providing a novel perspective of leveraging reviews as an external information source to supplement the visual details.

Besides inadequate descriptions of images, the online shopping experience of BLV people is also hindered by the weaknesses of screen readers dealing with crowded websites. Screen readers provide fine-grained information navigation and control, but at the cost of reduced walk-up-and-use convenience [34]. Prior work augmented the auditory web browsing experience by adding a secondary voice output [28, 42], supporting basic queries about the information provided by sellers [31], and supporting users to choose how many and which levels of detail to listen to based on their interest [23]. voice assistants as an alternative solution lack the ability to deeply engage with content and to get a quick overview of the landscape (e.g., list alternative search results & suggestions) [34]. Current voice assistants for online shopping such as *Alexa* only respond to a limited range of queries, e.g., *"add bananas to the cart"*, rather than finding targeted information to answer a BLV user's question. Inspired by VERSE [34] which extends a voice assistant with screen-reader-inspired capabilities to enhance web search, our work integrates review-driven question-answering with two levels of details (summary & original review lists), which contains rich first-hand insights from other buyers.

## 2.2 Information Retrieval In Online Reviews

Most image captioning and visual question answering benchmarks to date focus on questions such as simple counting and object detection that are directly based on the images. Marino *et al.* [19] draws on external knowledge resources when the image content is insufficient to answer customers' questions, but their focus on images containing general scenes cannot be directly leveraged for online shopping. Online reviews have been mentioned as an useful resource to assist BLV customers' desire for more product information [16, 30, 31], such as supporting question-answering [20, 36] and generating summaries [12, 15]. Prior work towards general users investigated extracting experiences [21, 24], tips [11, 41] or snippets suitable for product descriptions [7, 25], yet it remains

unknown how reviews can help with the unique interests of BLV users, especially providing descriptive information on product appearance. Our research fills a gap in the literature by leveraging existing human-authored resources of online reviews as an additional source of information to address such unique needs.

The most related work [25] applies a supervised method to extract reviews for image description based on 25K labeled sentences. However, the descriptions generated by this work are generic and rarely covered descriptions of visual details. Rather than following supervised methods such as conducting aspect-based sentiment analysis utilizing feature engineering [14, 35] or Bidirectional Encoder Representations from Transformers (BERT) [32] that requires a huge amount of human-labelled data, we proposed empirically-formulated syntactic rules based on our studies with BLV participants to retrieve meaningful review snippets for understanding appearances. Compared to "black box" data-driven approaches, these syntactic rules are explainable and can be validated by our targeted users; they are modular—existing rules can be edited or removed and new rules can be added without affecting the others. To the best of our knowledge, there is no prior work that directly informs rule design on extracting visual descriptions from reviews.

In summary, Revamp extends prior work on enhancing accessible information seeking by proposing a solution to bridge customer reviews with the needs of BLV users on visual information for a wider range of online products, including three frequently shopped categories: Home & Kitchen, Clothing, Shoes & Jewelry, and Electronics.

## 3 OVERVIEW OF STUDIES

In this work, we conducted three groups of studies:

- **Formative study (Section 4, 5)** with two stages to: (i) understand the current practice and derive design implications via 30-mins semi-structured interviews with 20 participants (P1-P20); (ii) investigate what specific visual information BLV users expect and how reviews can help via questionnaires and 30-mins semi-structured interviews with 10 participants (P11-P20).
- **Rule evaluation (Section 7.1)** with two stages to: (i) evaluate the informativeness of reviews extracted by the rules via questionnaires with eight participants (P13-P20); (ii) evaluate the rules' generalizability on different products where two sighted researchers rated the quality of the generated alt-text and answers in three informative levels.
- **System evaluation (Section 7.2)** to evaluate how the Revamp prototype affected the shopping process. Eight participants (P13-P20) browsed products with Amazon web pages and Revamp then provided subjective comments by online interviews. The study lasted for 40 mins.

The time interval between each group of study was about one month. Twenty participants (Gender: 7 female and 13 male; Age: average = 27.7, SD = 5.5) came from two different cultures (P1-P10: China; P11-P20: North America) participated in the first-stage formative study due to the distributed nature of the research team. Although the platforms and languages they use differ, participants share similar screen reader experience and information seeking challenges since the screen readers follow the general design and

<sup>1</sup><https://www.microsoft.com/en-us/ai/seeing-ai>

most online shopping websites are similarly structured with overloaded information. Only North American users were involved at the following studies because of the limited availability of review data on the Chinese shopping sites. Participants were compensated with \$15 USD/hour. Each study was audio recorded, transcribed and coded by three of the authors following the reflexive thematic analysis methods from Braun and Clarke [4]. The demographic information of participants is attached as Appendix.

## 4 UNDERSTANDING CURRENT PRACTICE AND DERIVING DESIGN IMPLICATIONS

We conducted 30-minute semi-structured online interviews with 20 BLV consumers (P1-P20) recruited through social media platforms. Specifically, the goal of this study is to understand that, when shopping online:

- What challenges do BLV users face in information seeking and what are their coping strategies?
- What are the design heuristics if we want to build an accessible information retrieval system for BLV users?

### 4.1 Challenges

Aligned with prior work [30, 31], two main challenges mentioned by our participants are lack of visual information and lack of efficient ways to navigate through information. We revealed how the two challenges were intertwined with each other, which elicited insights on designing an efficient way to leverage reviews to fill in the lacking visual information.

**4.1.1 Lack of Visual Information.** The most frequently mentioned challenge is the lack of detailed information for understanding the visual appearance of the product (18 out of 20 participants). While sighted users can tell the visual attributes such as color, shape and size or even functionalities of a product from a single image, the alt-texts of many current product images are either empty, or set as the image path, which does not provide any visual information for BLV users at all. Sellers usually provide no textual descriptions equivalent to the visual information conveyed by images, whether it is some basic attribute such as color and shape, or some vivid details such as pattern design. Take the color attribute as an example: the names provided by sellers sometimes can be too generic or obscure to interpret accurately. So much of the language we use to describe a product is centered around visual, e.g., “*marble pattern*” or “*Arctic blue*”, which poses a barrier to people who don’t experience the world visually. P18 brought up an example that a T-shirt with the color name “*surf the web*” was actually a kind of “*bright blue*”, which was only mentioned by a review.

**4.1.2 Lack of efficient ways to navigate through information.** Comparing with the scanning and skimming experience of sighted users, the passive and sequential reading manner of screen readers makes it inefficient for BLV users to retrieve useful information scattering all over a web page. P16 mentioned that it can be frustrating when accidentally jumping into another unrelated web component such as the shopping cart in the process of exploring product details. All participants agreed that reviews were useful resources. Review snippets with detailed descriptions can also provide further insights for product appearance. However, most participants (16 out of 20)

reported their utilization of reviews was greatly reduced because sifting through the large amount and unrelated reviews can be especially laborious. P19 stated that “*Usually there are too few reviews per page and there are no easy ways to jump from review to review. I know there exist useful reviews but it’s hard to directly pull them out without efforts.*”.

### 4.2 Current Solutions

There still exists a gap between current solutions and BLV users’ expectations of online shopping experience. Relying on human helpers, the same as many other scenarios, suffer from the limitation of availability, raises privacy concerns, and reduces BLV users’ independence. Meanwhile, as reviewed earlier in Related Work, many existing techniques lack adaptation to the specific needs of BLV users’ online shopping.

**4.2.1 Seeking Help from Sighted People.** 80% of the participants chose to seek help from sighted people including families/friends, other customers, paid video chat and accessibility hotlines. In this way, they had to trust the helper’s personal judgment and subject to issues such as helpers’ inavailability and privacy concerns. P19 mentioned that “*The shipping and delivery process is not the most time-consuming for me. Waiting for someone to help me with selecting products is.*” P12 stated that he did not always want others to know about certain products he buys and hoped there was an alternative solution to shop more independently.

**4.2.2 Leveraging Existing Web Functions and Automatic Tools.** P18 mentioned using visual interpretation apps such as Seeing AI to get brief generic descriptions on product images based on object and color detection, e.g., “*probably a close up of a red chair*”. Two participants who were more tech-savvy in web interfaces (P11, P20) mentioned using keyword-based search and filters, but these tools are under-utilized by other participants and can only remove irrelevant information on a general basis. P8 and P13 mentioned the customer questions & answers section supported by websites can also help with some generic details, yet it seldom covers appearance-related information, which is assumed available to sighted consumers via product images.

### 4.3 Implications for Design

Based on the interview results, we derived 3 design implications to improve the accessibility of online shopping experience. Specifically, the designs focus on leveraging the existing resources on a product page: *simplifying the webpage* (4.3.1) and *responding to active queries* (4.3.3) are to address the *lack of efficient ways to navigate through information* (4.1.2); *leveraging reviews to supplement visual description* (4.3.2) is to address the *lack of visual information* (4.1.1).

**4.3.1 Simplifying and reorganizing the webpage structure.** When the users’ current task is understanding a product, especially the appearance, the overloaded information and web components can become burdens for seeking visual-related description. To address this problem, we can provide an alternate view of the original page with less related information and components removed, such as shopping cart. Meanwhile, it is important to provide users the option to switch back to the original page.

**4.3.2 Leveraging customer reviews for extra visual description.** Participants' responses suggest that some subjective information in reviews can be helpful for BLV people to understand the visual attributes. Just as P2 stated: *"It is impossible for me to build a visual concept as you sighted people do, but I can feel what you feel. Reviews talking about the feelings when seeing the appearance are helpful for me to understand this pattern."* It remains to be investigated how to retrieve the informative reviews to meet the specific needs of BLV users.

**4.3.3 Responding to active queries with the page resources.** Participants expressed preferences to actively asking questions and getting summative answers of a product, which was usually what they did when consulting with a sighted helper. Meanwhile, any additional assistance mechanisms should be compatible with users' current screen reader experience to support users to access the original information as well. P6 commented *"I do not completely trust in answers selected by machines in case it may not cover very well"*.

## 5 LEVERAGING REVIEWS TO ANSWER BLV CUSTOMERS' QUESTIONS

The first-stage formative study uncovered the opportunities and challenges to support the unique needs of BLV population in visual information seeking in online shopping. Building upon these findings, we then seek to further understand:

- What specific visual or other product-related questions are BLV users interested in?
- How can we retrieve useful review snippets to answer these questions?

It is beyond the scope of this paper to exhaustively cover the large number of product categories and all the consumer questions corresponding to each category. To narrow down our focus, we first distributed questionnaires to ten North American participants (P11 - P20) who mainly use Amazon to collect their most frequently shopped product categories, then conducted one-on-one interviews to formulate the understanding of their specific questions on three main categories: Home & Kitchen, Clothing, Shoes & Jewelry, and Electronics, based on which we derived rule-based solutions to extract informative review snippets.

### 5.1 Product Categories and the Corresponding Question Types

The frequently shopped categories on Amazon by our participants are: Electronics (chosen by 90% participants), Home & Kitchen (50%), Pet Supplies (50%), Audible Books and CDs (50%), Clothing, Shoes & Jewelry (30%), Grocery and Gourmet Food (30%). Among these categories, we focused on three main categories whose appearances are as crucial as other information for making purchase decisions, including Home & Kitchen, Clothing, Shoes & Jewelry, and Electronics.

We then gathered 168 questions from participants in these three categories. Each participant was shown representative products on the Amazon best-sellers list and was asked to raise as many questions they want to ask as possible after browsing the title and description provided by sellers. The questions were then labeled and

divided into two groups, non-visual questions and visual questions. The distinct breakout of visual (57%) and non-visual (43%) questions is based on whether a question can be answered by directly observing the image.

First, 72 out of the 168 questions were questions on non-visual information that shares common interests with sighted consumers, including (i) functionality and other non-visual attributes, e.g., material, price; and (ii) consumer feedback, e.g., whether the product was worth the price, or of high quality.

We are interested in the remaining 96 questions on visual information, which were of specific interests to BLV users, including (i) visual attributes of image, e.g., color, logo, shape, size; and (ii) high-level concepts that can be inferred from an image, e.g., usage method, style.

### 5.2 Types of Visual Questions and How to Find Informative Answers

For non-visual questions, all participants reported that the useful information usually could be found in the product description provided by the sellers, reviews or customer questions & answers, while for visual questions, they mainly rely on human helpers to get the answers. Thus we focus our investigation on the following sub-categories of visual questions and whether we can provide informative answers to these visual questions based on reviews, hence providing an alternative for BLV users to shop online independently.

**5.2.1 Visual questions on basic attributes.** Most participants started with general appearance questions, such as *"how does it look like?"*, followed by detailed questions around the basic visual attributes of the product, including color, shape, size and logo, which puts forward the need of providing a briefing of appearance covering these basic attributes first, then responding to specific queries. Participants mentioned that they are not satisfied with only knowing the basic category name of visual attributes, or just vague comments, e.g., *"Nice shape. This color looks great"*. They feel interested in the specific detailed descriptions and impressions on appearance. P5 mentioned his story that once he wanted to buy a guitar and was curious about what exactly was the color attribute 'gold'. *"Is it a shining one or a matte one? I want to know more details."* For shape and size, participants always want to know how to compare the product with daily objects, e.g., *"Does it fit on a tabletop?"*, *"Is it like the size of a banana?"*, which can help them better understand shape and size in the real world. Participants are also curious about if there are logos on the product.

To provide informative answers on basic attributes mentioned above, the key is to extract the descriptive and comparative details on appearance from the large amount of reviews. Particularly, clothes can have more nuanced design details to describe, such as sleeves, waist, neck, edition of clothing, graphic designs or patterns. These questions, e.g., *"Do the buttons go all the way down?"*, *"What kind of neck design does the shirt have?"*, are very category-specific and often require fashion knowledge to provide aesthetic descriptions, as P14 mentioned *"even my friends don't describe well"*, hence descriptions in common customer reviews hardly meet their expectations. In this paper, we don't focus on these subdivided fashion knowledge which requires large amount of annotated data

(e.g. DeepFashion [17]), but focus on the common attributes shared by a wider range of products and possibly covered by customer reviews. The four attributes (color, shape, size, logo) are common across many products and can be concisely described in the alt-text compared to other aesthetic attributes (e.g., pattern, style).

**5.2.2 Visual questions on high-level concepts.** There are also high-level questions not directly related to basic visual attributes but can be inferred from the appearances in the image with commonsense, including usage, style, quality, texture and specific accessibility requirements, e.g., “Does it look like it has good quality? How to open this bottle? Does the product have physical buttons?” Such properties are similar to the “signifier” concept proposed by Don Norman<sup>3</sup>, which act as the indicator that can be interpreted meaningfully in the context of the social and physical world. These questions, usually framed as Visual Commonsense Reasoning, has been a challenging research topic in computer vision [40]. Yet currently it is not well-explored to address the need of inferring concepts and functions from product images for online shopping scenarios, as they require a large amount of annotated data from external knowledge sources beyond the shopping websites. However, the reviews commenting on the specific aspects (e.g. quality, physical buttons) can still indirectly provide insights on customers’ opinions. Thus our expediency is treating these questions same as the non-visual questions to provide the relevant reviews.

### 5.3 Rules-Based Solution for Retrieving Informative Review Snippets

In this work, we mainly focused on addressing BLV users’ questions on the common basic visual attributes, including color, logo, shape and size, which can be supported by augmenting keywords-based review searching with rule-based filtering (described below) to extract more informative snippets.

Compared to data-driven models, rule-based solutions can work well without the need of manually labelled data. Further, rules are explainable and can be validated by our targeted users. Rules are also modular—existing rules can be edited or removed and new rules can be added without affecting the others. Note that currently our rules are not designed to answer questions related to high-level concepts, which we consider as future work.

**5.3.1 Rules for Filtering Out Less Informative Reviews.** We first established several simple rules to *filter out* review sentences that cannot be used to further the understanding of visual attributes with specific rules:

(i) Short: sentence of 3 words or fewer; e.g., “satisfied”, “like the shape”, “poor logo”;

(ii) Reference to the image: reviews that comment whether the actual product is consistent with the photo provided by the seller, e.g., “looks exactly as pictured” “not shown as picture”. These sentences mentioned by our participants as “useless” and “annoying” since they often occur in the reviews but contain little useful information for BLV customers.

**5.3.2 Rules for Basic Visual Attributes.** All reviews containing the query keywords are included as candidates and annotated using

basic Part-Of-Speech tags<sup>4</sup>, which are also known as word classes or lexical categories, e.g., noun, adjective. Based on the aforementioned analysis, a group of three researchers iteratively established several rules for extracting informative reviews as follows.

**Rule 1: Adjective + Keyword or Keyword + Verb + Adjective.** The descriptive adjectives usually provide supplementary visual information. Examples include “a shimmery purple”, “crescent shape”, “a very nice-looking etched logo”, “suarish shape”, “The bubblegum color is glossy and fun.”, etc. It is hard to enumerate all the possible descriptive words, instead we decided to obtain qualified snippets by differentiating those with evaluative adjectives.

The evaluative adjectives expressing subjective emotions can be vague hence not helpful to further understand the visual attributes. e.g., “nice shape”, “great color”. We collected a list of such vague adjectives, including “great”, “nice”, good, “bad”, “horrible”, “disappointed” etc. These snippets are less informative but still provide subjective opinions, so we still keep them as candidates, but as the lowest priority among reviews following our rules.

**Rule 2: 1st pronoun + ... + Keyword + ... + that/which/but/because.** Rather than the simple sentences only containing expressions on attitude e.g., “I feel disappointed at the color.”, the sentences with clauses usually provide more detailed and useful information. Coordinating conjunction, e.g., with “but” emphasizes the two statements contrasting or contradicting each other. Subordinating conjunctions, e.g., with “because” often contain detailed reasons for customers’ attitude towards the visual attribute. Examples include “I love the color (Bubblegum), which I bought because it was the lowest cost for a color that would be difficult to misplace or forget while traveling”.

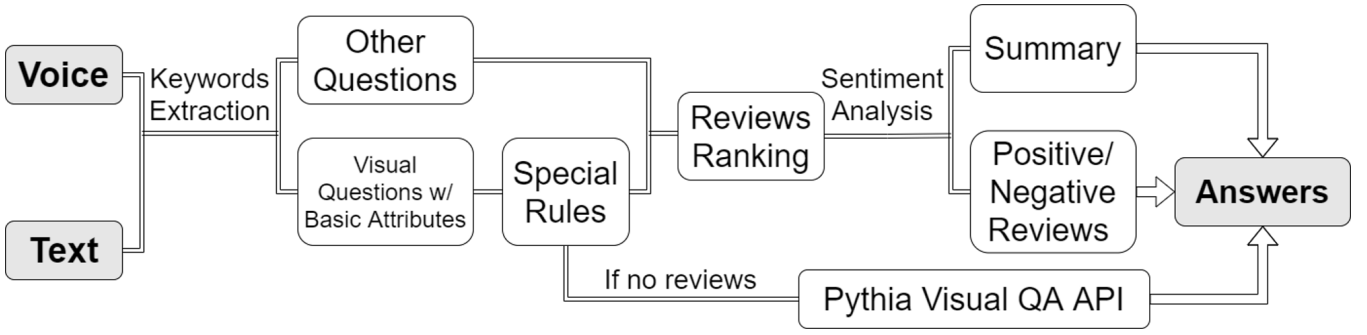
**Rule 3: Comparative Expressions.** Other than the common rules above, there exist some specific rules that work for particular visual attributes. (1) Keyword (shape) + “like/liked”, e.g., “shaped like a Cola Bottle”. Comparing the shape of a product with a familiar daily object can be helpful for better understanding the shape; (2) Keyword (size) + “fit/for/of”, e.g., “size fits in all cup holders”, shows reviews containing details on how the product fits in the settings are informative; (3) “than/more of” + Keyword (color), e.g., “it is a terra cotta than mocha”. Some products look different with the picture provided by the sellers. Sighted customers complaining about this kind of difference can also be informative for better understanding the actual appearance.

### 5.4 Generating Answers and Image Descriptions

Finally, we use the review snippets retrieved by the rules above to answer a BLV user’s specific visual question or to compose a brief appearance description. As an answer to a BLV user’s question, we provide original reviews retrieved by our rules divided into two lists—positive and negative—based on a sentiment analysis. We also provide a summary with the numbers of positive and negative reviews, and the top-3 informative review snippets from reviews across the two lists. To generate both the lists and the top-3 summary, we need to rank the review snippets. In particular, review snippets selected by Rule 1 with descriptive adjs, 2, and 3 have higher priority than those by Rule 1 with evaluative adjs; further,

<sup>3</sup><https://sites.google.com/site/thedesignofeverydaythings/home/signifiers>

<sup>4</sup><http://www.nltk.org/book/ch05.html>



**Figure 2: System overview of how Revamp responds to users' query. It first extract the keywords and process the query by its kind. After filtering out irrelevant reviews and reranking, Revamp uses the shortest candidate among the top-3 reviews to generate the visual description of a product and relevant answers to the query.**

reviews of the same priority are ranked by their helpfulness (votes by other customers). The next level lower are reviews that contain the query keywords but do not meet any rules, within which they are ranked by relevance based on the concept of graph centrality following prior work [37]. For the brief appearance description, we select the shortest sentence in the top-3 candidates corresponding to each visual attribute (where each visual attribute is used as a keyword) since it is advised not to use long text for image alt-text. If users are interested in learning more details, they can move forward to ask specific questions on visual attributes.

## 6 IMPLEMENTATION OF REVAMP—AN INFORMATION RETRIEVAL SYSTEM FOR ONLINE SHOPPING

We present Revamp, an interactive information retrieval system to improve the shopping experience for BLV users. We implemented Revamp as a browser extension that works on Amazon. Revamp allows BLV users to interact with a simplified product page, access image description composed of relevant reviews, ask questions and receive a summary with original reviews as response, as shown in Figure 2. On the front-end, Revamp uses Chrome API to automatically simplify the page and injects our search component to the current page, allowing user to send query and get a response from our back-end. On the back-end, Revamp extracts the keywords in user's query and matches the keywords with review snippets data. By ranking the filtered review snippets and applying sentimental classification, the back-end returns a review summary and a positive and a negative lists.

**Data source.** Our data source consists of two main parts: (i) Basic attributes (title, color, price) from information provided by sellers; (ii) Reviews (containing title, content, rating, helpfulness, date, and author); and (iii) Customer Q&A; We scraped data from Amazon product pages using the Python library BeautifulSoup combined with Selenium and stored it in .csv format. If a product browsed by a user already exists in our database, we will use it directly; otherwise, we will run the scraper to retrieve the data and save it in the database. Our Python Flask based back-end API server will use these data to generate responses to our front-end's requests.

Before our user study and evaluation, we pre-scraped the product data to avoid the potential latency of retrieving it in real-time.

**Web page simplification** We provide an alternate view of the original Amazon pages to improve the screen reader browsing experience. Our extension firstly rearranges the elements on the Amazon product page by manipulating the DOM tree. We removed the irrelevant information such as advertisements and promotions, and keep only the product details, reviews and images from the original page. Revamp would generate brief descriptions as the alt-text of the images using the aforementioned rules. Then we use Chrome APIs to inject our Revamp module into the product page. We strictly followed WAI-ARIA<sup>5</sup> standards to make sure all components have proper attributes and keyboard interactions.

**Review Snippets Extraction.** Revamp responds to both keyword queries such as “color” as well as natural language questions such as “Does the product have physical buttons?”. For keyword queries on four main visual attributes, including color, logo, shape, size, we extracted the reviews with the pre-defined extended keywords based on our formulated rules. Specifically, for color, the extended keywords included ‘color’, basic color name such as ‘blue’ and the special name provided by sellers such as ‘surf the web’. For other natural language questions, we used the Rapid Automatic Keyword Extraction (RAKE) library<sup>6</sup> for keyword extraction, Synset from WordNet<sup>7</sup> to get the groupings of synonymous keywords that express the same concept, then extracted the reviews with the certain and synonymous keywords.

**Answer Generation.** There are three types of information generated. (i) A supplementary textual description of product images; (ii) Two reviews lists: positive and negative; and (iii) A summary of reviews based on the two reviews lists. After we retrieved the ranked reviews, we divided them into positive and negative categories based on aspect-based sentimental classification by MonkeyLearn API. Finally, we extract top three reviews from each list to generate our summary using a pre-defined template. In case there are too few reviews to populate the answers, we leverage the

<sup>5</sup>WAI-ARIA stands for Web Accessibility Initiative – Accessible Rich Internet Applications. It is a specification published by the World Wide Web Consortium (W3C) that specifies how to increase the accessibility of web pages.

<sup>6</sup><https://github.com/csurfer/rake-nltk>

<sup>7</sup><https://www.nltk.org/howto/wordnet.html>



The product you are shopping for is a T-shirt with a color name as "Surf the web". Which of the following review snippets do you think would be helpful for understanding the product color better? \*

- ☐ We ordered 2 blue and black in size xl. The blue was ok and the black was small
- ☐ Wasn't a royal blue color was very disappointed in it
- ☐ Bought the "surf the web" blue and the color is just as vibrant as the picture.
- ☐ I prefer a slightly darker blue but it's cool.
- ☐ color was off and not the true blue/normal blue that champion usually has
- ☐ I bought the blue one and was thinking about getting other colors

**Figure 3: An example questionnaire for BLV users to vote for which one of the review snippets (retrieved by our rules or by default ranking of search result) best answers visual attributes questions.**

existing computer vision technique to provide basic answers, e.g., using Pythia [13] to answer the color, whether there is a logo on the image, and what is the shape. Users can either use text or voice input, when Revamp finishes generating answers, it will notify the user with a beep. Rather than automatically reading out the answers, we support users to use their own screen readers to scan the answers, which provides them with more freedom of control.

## 7 EVALUATION

We separated evaluation into two parts: rule evaluation and system evaluation. Rule evaluation is to evaluate the performance of rules-based solution for retrieving informative review snippets. System evaluation is to further evaluate how Revamp is integrated into BLV users' online shopping process. Modifying page structure and reviews-based interactions are non-separable, as they are coordinated to support a comprehensive information seeking flow. Hence we chose to evaluate Revamp as an integrated system.

### 7.1 Rule Evaluation

We first evaluated how our rules cover the informative review snippets for understanding appearance and how our solution can be generalized to different products.

**7.1.1 Coverage of Informative Reviews for Understanding Appearances.** We randomly selected eight products covering three aforementioned main categories from the Amazon best-seller list and extracted review snippets covering four main visual attributes (color, logo, shape and size) following the above rules. We then distributed voting questionnaires to eight BLV participants (P13 - P20). Given each product, a participant was asked to vote for the snippet that could help them best understand the visual attributes of that product: half of the snippets were generated by our proposed rules and the other half from the top-ranked answers in the existing Amazon 'Have a Question' section; all snippets were presented in a random order. Figure 3 shows an example question and review snippets. The snippets selected by our rules gained 85% of the 144 votes. Notice that we are not filtering out the remaining 15%. Review snippets

which were not selected by our rules may also contain informative details, but seldom directly for describing the visual attributes, e.g., a review "We ordered 2 blue and black in size xl. The blue was ok and the black was small" talking about color actually gave more details about the size differences between different colors. Rather than filtering out these reviews, we still show them but with lower ranking, while reviews that meet our rules and provide descriptive details such as "color was off and not the true blue/normal blue that champion usually has" will be assigned a higher priority.

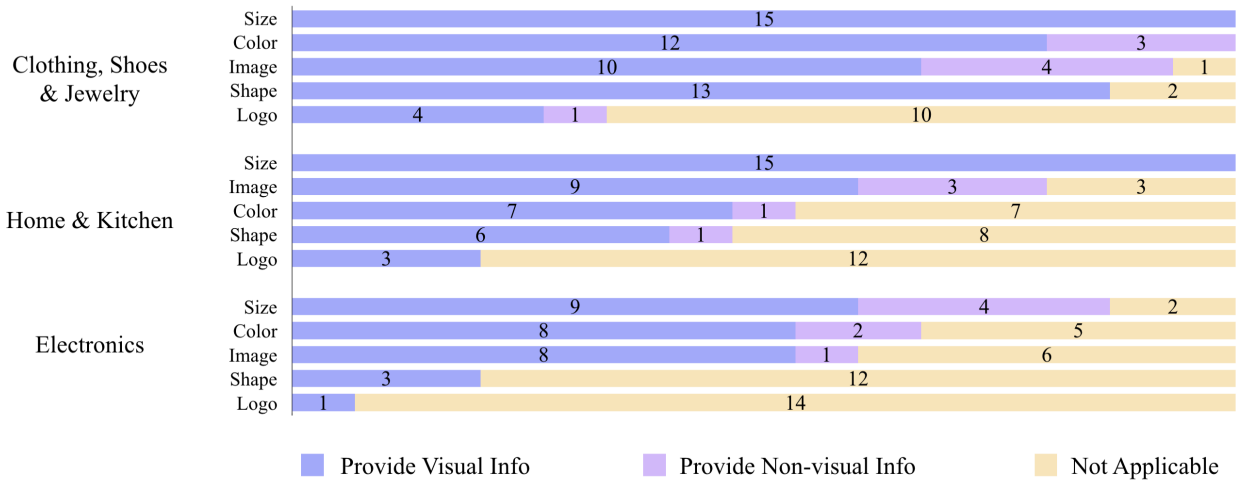
**7.1.2 Generalization to Various Products.** To further validate the generalization of these rules in a more real-world setting, we downloaded web data of 45 selected products from Amazon's up-to-date top-seller ranking list (top 15 in each of the aforementioned three categories with similar products removed to maintain variety). Two sighted research team members then browsed the product page and rated the quality of the generated alt-text and answers that describe four visual (color, logo, shape, size) in three levels: *not applicable*, *providing related non-visual information*, *providing direct visual information*. Results showed that our rules can provide informative results in most cases as long as there exist enough reviews, especially for the products whose appearances matter more in shopping decision, as shown in Figure 4. Not applicable is due to no reviews mentioned the attribute if the specific attribute does not occur (there is no logo on the product), or has no details to provide (e.g., no reviews of a white bed sheet mentioned color details), or there exist too few informative reviews. In this case, Revamp uses Pythia Visual Question Asking [13] to provide back-up answers.

### 7.2 System Evaluation

To further understand how the system is integrated into BLV users' online shopping process, we evaluated Revamp with BLV consumers. We selected six representative products from the aforementioned three main categories including Water Bottle, Men's Tee, Women's Dress, Bed Sheet, Bluetooth Speaker, and Chair Cushion as the testing products. Three representative products, retrieved reviews by rules and generated image description are shown in Figure 5.

**7.2.1 Participants, Design, and Procedure.** We conducted interviews with the same eight blind participants (P13 - P20). Each user study lasted approximately 40 mins using Zoom after we received the IRB approval. Participants used their own laptops and screen readers and shared their feedback with the experimenters at the same time. We first asked the users to browse the products as they used to on original Amazon web pages without Revamp. The main task was to obtain information on four visual attributes (color, logo, shape, size) also other details they were curious about. We then introduced how Revamp worked, including how it simplified the original Amazon web page and what kinds of questions Revamp could answer, and asked the participants to repeat the same task using Revamp. Participants then filled out a survey with Likert-scale statements as shown in Table 1. Participants then further explained their reasons of giving each specific ratings. Each study was audio recorded and transcribed and participants' qualitative response was summarized by affinity diagramming.





**Figure 4: Performance of image description and visual questions on four basic attributes including shape, logo, size, color. Two sighted research team members rated the quality by three levels on 15 Amazon top-seller products in each of the three main categories. Results showed that our rules can directly provide visual info for 54 questions or image descriptions in Clothing, Shoes & Jewelry; 40 in Home & Kitchen; 29 in Electronics. We also included “Not Applicable” because of no relevant reviews or too few informative reviews.**

**Table 1: We collected participants’ subjective ratings on Revamp. The scale was 1 to 7, in which 7 = I strongly agree with this statement, 4 = neutral, and 1 = I strongly disagree with this statement. The value represents the average of ratings (SD). Data was analyzed using Wilcoxon test and a statistical significant difference ( $p < 0.05$ ) is marked with \*—all ratings of Revamp significantly outperformed the current practice.**




Statements	Revamp	Current
It is easy and efficient for me to locate product-related information I need *	5.50 (1.07)	4.75 (0.89)
My questions are answered with informative answers *	5.13 (0.64)	4.38 (1.41)
I feel confident in understanding the appearance of the product *	5.88 (1.25)	3.00 (0.76)
I believe I can fully utilize the information from reviews *	5.75 (1.83)	5.13 (1.46)
Revamp can be a helpful alternative for shopping online independently when no sighted helper is available	6.50 (0.53)	-
I will use Revamp regularly in my daily life	6.00 (0.76)	-

**7.2.2 Results.** We observed the common interaction pattern with Revamp proceeded as follows: The participants first navigated through four elements (product name, image with alt-text, info provided by sellers and QA) on Revamp. They then asked questions on visual attributes and other aspects they were curious about (no predefined questions were provided). Since personal interests differ, they asked 2 ~5 questions for each product, including both visual and non-visual questions. Besides the review summary, they also browsed the positive and negative review lists when needed.

Overall participants considered Revamp to be a helpful alternative for shopping online independently when no sighted helper is available and will use Revamp regularly in daily life. The comparative ratings are shown in Figure 6. Notice that one limitation of

study is the bias of comparing a new system to an old one, participants would know which system was designed by the authors and can be influenced by this. Unfortunately, given the popularity of Amazon, participants would have known which one was ours even if we had intentionally de-identified the two systems. We further discuss participants’ subjective comments below.

**Providing supplementary information for better understanding the product appearance.** With Revamp, participants’ ratings on their understanding of product appearance were improved comparing with their current practice, owing to the supplementary image description and responses to visual questions. Having a supplementary description of the product image is a new

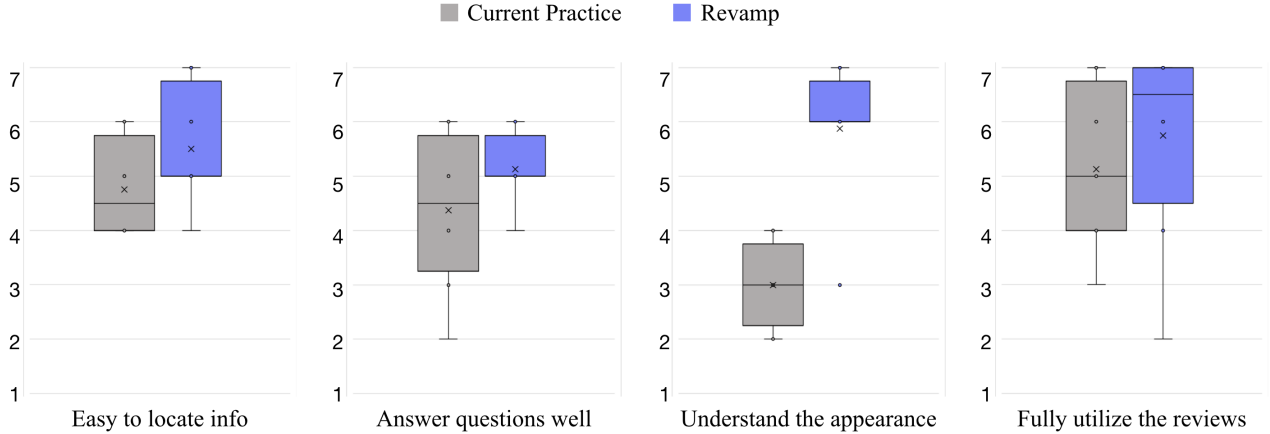
Category	Rule	Retrieved Reviews	Image Description by Revamp	Recognition by SeeingAI
Clothing, Shoes & Jewelry 	Rule 1	I believe is called "army green" is actually <b>more of a gorgeous kelly green</b> that is perfect for Christmas, <b>Green is vibrant</b> and true to color	Green is vibrant and true to color, it was a little loose but overall I'm glad I went a size up, Cute and flattering shape, <u>no logo (from Pythia API)</u>	Probably a green dress
	Rule 2	It was a little loose <b>but</b> overall I'm glad I went a <b>size</b> up		
	Rule 1	Cute and <b>flattering shape</b> , Perfect for <b>hour glass shaped</b> ladies		
Home & Kitchen 	Rule 3	<b>Spice</b> is like a <b>burnt orange</b> , if you ask me it's more salmon <b>than spice</b>	Spice is like a burnt orange, Oval shaped, It's a giant size cushion, <u>no logo (from Pythia API)</u>	Probably a person sitting on a leather couch
	Rule 1	<b>Oval shaped</b> , cushion <b>shape</b> is not <b>round</b> , <b>curvy (Concave) shape</b>		
	Rule 1, Rule 3	It's a <b>giant size</b> cushion, <b>fits well for</b> the <b>size</b> compared to the frame (44" frame) and it is very plush		
Electronics 	Rule 1, Rule 2	The <b>color</b> is <b>beautiful</b> , a <b>light aquamarine</b> that is very <b>mild</b> and <b>unoffensive</b> . the <b>aquamarine color</b> is <b>playful and fun</b> .	the aquamarine color is playful and fun, almost about the size of a brick, has a clean look to it with its squarish shape, <u>has logo (from Pythia API)</u>	Probably a close up of electronics
	Rule 1	has a clean look to it with its <b>squarish shape</b>		
	Rule 3	almost about the <b>size of a brick</b> , Is the perfect <b>size for</b> so many uses		

**Figure 5: Examples of representative products in three main categories, showing the retrieved reviews, generated image description and comparing with the recognition results by SeeingAI, a common tool for BLV users to figure out what is in an image.**

experience to them: “You know the Amazon doesn’t provide image descriptions for the products. I will definitely use this add-on in my life.” (P13) When using Amazon without Revamp, participants often just quickly skipped the image web component since no informative information is provided; with Revamp, participants tended to use the shortcut command of screen readers to directly jump to image to first gain an overall understanding of the product’s appearance.

Participants felt surprised when Revamp provides useful answers on visual attributes, as mentioned by P19: “It is really cool. It did answer my questions about product appearance.” and even helps them learn about unfamiliar visual attributes, as shared by P13: “Although I haven’t seen colors before, I have a lot of fun in reading these descriptions of colors. For example, I don’t know what is the ‘Spice’ color. It told me a review mentioned ‘like a burnt orange’, which is much more understandable.” Participants also mentioned that descriptive details of reviews from the first-hand buyers sometimes felt more trustworthy than opinions of a friend or family member, as P15 commented “Who can perform better than the customers who have bought the product themselves on describing the product? I really like the idea of using the reviews.” Most participants agreed that they can better utilize the information from reviews. Only one participant gave a low rating of 2 because he personally preferred not to use any filter on the reviews in case something important is missed.

**Enhancing interaction flow of information seeking.** With Revamp, the information seeking experience of participants was improved comparing with their current practice, owing to the reconstructed web page and better utilizing reviews. Revamp provided users with cleaner web page structure and is more user-friendly than the current Amazon page, as stated by P20: “I don’t need to worry of being stuck in useless information any more.” Participants found the review summary and the review lists divided by positive and negative useful to access customers’ opinions more efficiently and proactively. P14 mentioned that “I like it that Revamp also keeps the original reviews accessible in the lists. After reading the summary, I can then make the decision to skip or look into the details of each relevant reviews in the list.” Participants were more engaged in asking questions, as P19 commented “Sometimes I could be inspired to ask more after I read the answers.” Participants could choose to interact with the system using either voice commands or text input based on their preferences. Although participants in general liked the experience of screen readers augmented with voice input, they preferred to receive answers in text rather than speech: “Using my screen reader to read out the answers is far more better than directly answering my questions by voice. I can adjust the speed and pause anytime.” (P15)



**Figure 6: Subjective ratings on statements comparing current practice with experience of using Revamp. Revamp demonstrates a clear advantage in the experience of understanding the product appearance.**

## 8 DISCUSSION

Most web pages with vivid designs and a large amount of information are still not accessible enough to BLV users at present. We explored three aspects towards enhancing the information seeking experience on online shopping websites: (i) simplifying and reconstructing the web pages according to users’ current task; (ii) providing coordinated experience of active query and passive reading to support flexible information seeking; (iii) leveraging relative text resources on the web page, such as reviews, to fill in the information gap. Besides, this work also inspired several exciting future directions as follows.

### 8.1 Working with Multiple Product Pages of Online Shopping

In this work, we focus on the task of understanding the product details, especially the appearance. Imagine another task: if the user has several product candidates in mind, we should also correspondingly meet a new information seeking need of comparing multiple products. To address this, future work can explore reconstructing the webpages to support efficient switching among products and answering comparative questions about multiple products. Currently, table is a common web page element for comparing different products and should be supported in future versions of Revamp. For example, the system should guide the user to navigate a table by informing them which products are being compared in terms what attributes and further be able to answer a user’s question by looking up the table.

### 8.2 Working with A Broader Variety of Product Pages

Furthermore, future work can employ and test our methods on other product pages where there is often an overload of information

inaccessible for BLV users to retrieve. For example, on Yelp, our methods can add description to popular dishes, whose images do not always convey all important information (e.g., how large is the portion); on Trip advisor, images taken by fellow travellers can be further described using our methods by extracting relevant descriptions from others’ comments (e.g., whether a trail has shade); on Youtube, comments can be leveraged to generated video captions for eye-catching moments, which serve as a more vivid introduction before one decides whether to watch a given video.

### 8.3 Working with Supervised methods

Although the subjective comments were useful and convey vivid details, our participants pointed out that the generated appearance descriptions might be too subjective since it only contained one snippet in each visual aspect. However, if we follow prior work on filtering out subjective comments [25], we will lose the vivid and descriptive details that are crucial to appearance understanding. As such, participants preferred the image description to remain as concise as possible and the description provided by Revamp would serve more as “a first step or a clue.” (P16) Also, it remains to be explored how we can better retrieve useful information from reviews on the high-level visual concepts inferred by image. In this work, we formulated hand-crafted rules to explore the possibility of leveraging reviews as an additional information source to fill the visual information gap. In the future, we can collect and label data to deploy supervised methods such as Knowledge Graph which could extract more information to better support queries not limited in the four visual aspects (color, logo, shape and size) and provide better summative appearance description.

### 8.4 Working with Better QA Experience

The question & answering experience of Revamp can be improved by providing prompting questions and involving more information

sources. We observed that some participants had difficulty in formulating what questions to ask: P16 only asked about basic attributes, e.g., color, quality. She explained that “*Sometimes I don’t know how to ask a ‘good’ question. Maybe if I change my wording, I can get more answers.*” and suggested us to provide more pre-defined questions as prompts. Meanwhile, participants noticed that there still exists a gap between their expectations and the answer quality in other details beyond the four main visual attributes. “*Revamp can support descriptive answers for basic visual questions, but when I want to ask about some concrete (non-visual) details such as the dimensions of the bottle, it cannot answer me directly.*” (P19) In the future, we can involve more information source into Revamp, also leveraging Optical Character Recognition to extract the text information on the images provided by sellers.

### 8.5 Working in the Real-World Settings

Similar to other information retrieval system, the performance of Revamp highly depends on the quality of existing data. Currently, we leverage existing visual question answering API to provide a back-up answer for visual questions when there are too few reviews or no informative reviews. In the future, we can also involve human helpers in the system by sending the questions to crowd workers. Our rules for extracting informative reviews can also give insights on formulating guidelines of describing products for BLV users: (i) include descriptive details rather than vague opinions (Rule 1); (ii) if have to express personal attitudes, elaborate the reasons (Rule 2); (iii) compare with daily or common objects to help understand new concepts (Rule 3). Besides, participants also hoped to see Revamp’s functions integrated into the mobile application with Alexa since they often try the mobile application of Amazon when faced with accessibility problems shopping on the web.

### 8.6 Limitations

There are several limitations of the system evaluation study. First, we did not counterbalance the conditions in the system evaluation. Using Revamp first would introduce a strong carry-over effect, as the information provided by Revamp is an augmentation of the baseline. As personal interests and product information differ, to compare intuitively how Revamp improves the experience, all participants first browsed the products as they used to on original Amazon web pages then on the modified web page of the same product by Revamp. Second, there exists the bias of comparing a new system to an old one when participants know which system was designed by the researchers [5, 33]. Third, only eight participants and a limited amount of products were involved in evaluation. This limitation can be addressed by a large-scale in-the-wild study, which we regard as future work beyond this initial study.

## 9 CONCLUSION

We present Revamp, a system that employs information retrieval techniques to meet the unique information seeking requirements of BLV consumers when shopping online. Our main contribution is a rule-based approach that leverages rich customer reviews to serve as image description and to answer BLV users’ questions related to product appearances. Evaluations with eight BLV consumers showed that Revamp provides useful subjective information for

understanding the product appearance and enhances the accessible information seeking experience of online shopping. Although Revamp could not provide answers for all of the products (e.g., when there are too few reviews of a product), it can serve as an effective supplemental helper for BLV users to better access and understand a product before making a purchase decision.

## ACKNOWLEDGMENTS

The authors thank all participants who generously shared their time and experience for this work, Mengqi Li for helping with formative studies, and Prof. Jacob O. Wobbrock for valuable advice on the paper writing.

## REFERENCES

- [1] Mauro Avila, Katrin Wolf, Anke Brock, and Niels Henze. 2016. Remote Assistance for Blind Users in Daily Life: A Survey About Be My Eyes. In *Proceedings of the 9th ACM International Conference on Pervasive Technologies Related to Assistive Environments* (Corfu, Island, Greece) (PETRA ’16). ACM, New York, NY, USA, Article 85, 2 pages. <https://doi.org/10.1145/2910674.2935839>
- [2] Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. 2010. VizWiz: Nearly Real-time Answers to Visual Questions. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) (UIST ’10). ACM, New York, NY, USA, 333–342. <https://doi.org/10.1145/1866029.1866080>
- [3] Erin Brady, Meredith Ringel Morris, Yu Zhong, Samuel White, and Jeffrey P. Bigham. 2013. Visual Challenges in the Everyday Lives of Blind People. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI ’13). ACM, New York, NY, USA, 2117–2126. <https://doi.org/10.1145/2470654.2481291>
- [4] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806> arXiv:<https://doi.org/10.1080/2159676X.2019.1628806>
- [5] Emeline Brulé, Brianna J. Tomlinson, Oussama Metatla, Christophe Jouffrais, and Marcos Serrano. 2020. Review of Quantitative Empirical Evaluations of Technology for People with Visual Impairments. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI ’20). Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376749>
- [6] Michele A. Burton, Erin Brady, Robin Brewer, Callie Neylan, Jeffrey P. Bigham, and Amy Hurst. 2012. Crowdsourcing Subjective Fashion Advice Using VizWiz: Challenges and Opportunities. In *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility* (Boulder, Colorado, USA) (ASSETS ’12). Association for Computing Machinery, New York, NY, USA, 135–142. <https://doi.org/10.1145/2384916.2384941>
- [7] Guy Elad, Ido Guy, Slava Novgorodov, Benny Kimelfeld, and Kira Radinsky. 2019. Learning to Generate Personalized Product Descriptions. 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, 389–398. <https://doi.org/10.1145/3357384.3357984>
- [8] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. Every Picture Tells a Story: Generating Sentences from Images. In *Proceedings of the 11th European Conference on Computer Vision: Part IV* (Heraklion, Crete, Greece) (ECCV’10). Springer-Verlag, Berlin, Heidelberg, 15–29.
- [9] Cole Gleason, Amy Pavel, Emma McCamey, Christina Low, Patrick Carrington, Kris M. Kitani, and Jeffrey P. Bigham. 2020. Twitter A11y: A Browser Extension to Make Twitter Images Accessible. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI ’20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376728>
- [10] Darren Guinness, Edward Cutrell, and Meredith Ringel Morris. 2018. Caption crawler: Enabling reusable alternative text descriptions using reverse image search. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–11.
- [11] Ido Guy, Avihai Mejer, Alexander Nus, and Fiana Raiber. 2017. Extracting and ranking travel tips from user-generated reviews. In *Proceedings of the 26th international conference on world wide web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 987–996.
- [12] Jeff Huang, Oren Etzioni, Luke Zettlemoyer, Kevin Clark, and Christian Lee. 2012. RevMiner: An Extractive Interface for Navigating Reviews on a Smartphone. In

- Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) (UIST '12). ACM, New York, NY, USA, 3–12. <https://doi.org/10.1145/2380116.2380120>
- [13] Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2018. Pythia v0.1: the Winning Entry to the VQA Challenge 2018. arXiv:1807.09956 [cs.CV]
  - [14] Svetlana Kiritchenko, Xiaodan Zhu, Colin Cherry, and Saif Mohammad. 2014. NRC-Canada-2014: Detecting aspects and sentiment in customer reviews. In *Proceedings of the 8th international workshop on semantic evaluation (SemEval 2014)*. Association for Computational Linguistics, Dublin, Ireland, 437–442.
  - [15] Fangtao Li, Chao Han, Minlie Huang, Xiaoyan Zhu, Ying-Ju Xia, Shu Zhang, and Hao Yu. 2010. Structure-aware review mining and summarization. In *Proceedings of the 23rd international conference on computational linguistics*. Association for Computational Linguistics, Association for Computational Linguistics, Stroudsburg, PA, 653–661.
  - [16] Guan hong Liu, Xianghua Ding, Chun Yu, Lan Gao, Xingyu Chi, and Yuanchun Shi. 2019. "I Bought This for Me to Look More Ordinary": A Study of Blind People Doing Online Shopping. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 372, 11 pages. <https://doi.org/10.1145/3290605.3300602>
  - [17] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, Long Beach, CA, 1096–1104.
  - [18] Christina Low, Emma McCamey, Cole Gleason, Patrick Carrington, Jeffrey P. Bigham, and Amy Pavel. 2019. Twitter A11y: A Browser Extension to Describe Images. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 551–553. <https://doi.org/10.1145/3308561.3354629>
  - [19] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Long Beach, CA, 3195–3204.
  - [20] Julian McAuley and Alex Yang. 2016. Addressing complex and subjective product-related queries with customer reviews. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 625–635.
  - [21] Hye-Jin Min and Jong C Park. 2012. Identifying helpful reviews based on customer's mentions about experiences. *Expert Systems with Applications* 39, 15 (2012), 11830–11838.
  - [22] Meredith Ringel Morris, Kori Inkpen, and Gina Venolia. 2014. Remote Shopping Advice: Enhancing In-store Shopping with Social Technologies. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Baltimore, Maryland, USA) (CSCW '14). ACM, New York, NY, USA, 662–673. <https://doi.org/10.1145/2531602.2531707>
  - [23] Meredith Ringel Morris, Jazette Johnson, Cynthia L. Bennett, and Edward Cutrell. 2018. Rich Representations of Visual Content for Screen Reader Users. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3173574.3173633>
  - [24] Quang Nguyen. 2012. Detecting experience revealing sentences in product reviews. *University of Amsterdam* 1, 1 (2012), 1–78.
  - [25] Slava Novgorodov, Ido Guy, Guy Elad, and Kira Radinsky. 2019. Generating Product Descriptions from User Reviews. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 1354–1364. <https://doi.org/10.1145/3308558.3313532>
  - [26] Krishnan Ramnath, Simon Baker, Lucy Vanderwende, Motaz El-Saban, Sudipta N Sinha, Anitha Kannan, Noran Hassan, Michel Galley, Yi Yang, Deva Ramanan, et al. 2014. Autocaption: Automatic caption generation for personal photos. In *IEEE Winter Conference on Applications of Computer Vision*. IEEE, Waikoloa Village, Hawaii, 1050–1057.
  - [27] André Rodrigues, Kyle Montague, Hugo Nicolau, João Guerreiro, and Tiago Guerreiro. 2017. In-context Q&A to Support Blind People Using Smartphones. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (Baltimore, Maryland, USA) (ASSETS '17). ACM, New York, NY, USA, 32–36. <https://doi.org/10.1145/3132525.3132555>
  - [28] Daisuke Sato, Shaojian Zhu, Masatomo Kobayashi, Hironobu Takagi, and Chieko Asakawa. 2011. Sasayaki: Augmented Voice Web Browsing Experience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). ACM, New York, NY, USA, 2769–2778. <https://doi.org/10.1145/1978942.1979353>
  - [29] Kristen Shinohara and Jacob O. Wobbrock. 2011. In the Shadow of Misperception: Assistive Technology Use and Social Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). ACM, New York, NY, USA, 705–714. <https://doi.org/10.1145/1978942.1979044>
  - [30] Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 2020. "Person, Shoes, Tree. Is the Person Naked?" What People with Vision Impairments Want in Image Descriptions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376404>
  - [31] Abigale J. Stangl, Esha Kothari, Suyog D. Jain, Tom Yeh, Kristen Grauman, and Danna Gurari. 2018. BrowseWithMe: An Online Clothes Shopping Assistant for People with Visual Impairments. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility* (Galway, Ireland) (ASSETS '18). ACM, New York, NY, USA, 107–118. <https://doi.org/10.1145/3234695.3236337>
  - [32] Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 380–385. <https://doi.org/10.18653/v1/N19-1035>
  - [33] Shari Trewin, Diogo Marques, and Tiago Guerreiro. 2015. Usage of Subjective Scales in Accessibility Research. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers & Accessibility* (Lisbon, Portugal) (ASSETS '15). Association for Computing Machinery, New York, NY, USA, 59–67. <https://doi.org/10.1145/2700648.2809867>
  - [34] Alexandra Vtyurina, Adam Fournery, Meredith Ringel Morris, Leah Findlater, and Ryen W. White. 2019. VERSE: Bridging Screen Readers and Voice Assistants for Enhanced Eyes-Free Web Search. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 414–426. <https://doi.org/10.1145/3308561.3353773>
  - [35] Joachim Wagner, Piyush Arora, Santiago Cortes, Utsab Barman, Dasha Bogdanova, Jennifer Foster, and Lamia Tounsi. 2014. DCU: Aspect-based Polarity Classification for SemEval Task 4. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*. Association for Computational Linguistics, Dublin, Ireland, 223–229. <https://doi.org/10.3115/v1/S14-2036>
  - [36] Mengting Wan and Julian McAuley. 2016. Modeling ambiguity, subjectivity, and diverging viewpoints in opinion question answering systems. In *2016 IEEE 16th international conference on data mining (ICDM)*. IEEE, IEEE, Barcelona, Spain, 489–498.
  - [37] Vinicius Woloszyn, Henrique D. P. dos Santos, Leandro Krug Wives, and Karin Becker. 2017. MRR: An Unsupervised Algorithm to Rank Reviews by Relevance. In *Proceedings of the International Conference on Web Intelligence (Leipzig, Germany) (WI '17)*. Association for Computing Machinery, New York, NY, USA, 877–883. <https://doi.org/10.1145/3106426.3106444>
  - [38] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. 2017. Automatic alt-text: Computer-generated image descriptions for blind users on a social network service. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, Association for Computing Machinery, New York, 1180–1192.
  - [39] Chien Wen Yuan, Benjamin V. Hanrahan, Sooyeon Lee, Mary Beth Rosson, and John M. Carroll. 2017. I Didn't Know That You Knew I Knew: Collaborative Shopping Practices Between People with Visual Impairment and People with Vision. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 118 (Dec. 2017), 18 pages. <https://doi.org/10.1145/3134753>
  - [40] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Long Beach, CA, 6720–6731.
  - [41] Di Zhu, Theodoros Lappas, and Juheng Zhang. 2018. Unsupervised tip-mining from customer reviews. *Decision Support Systems* 107 (2018), 116–124.
  - [42] Shaojian Zhu, Daisuke Sato, Hironobu Takagi, and Chieko Asakawa. 2010. Sasayaki: An Augmented Voice-based Web Browsing Experience. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility* (Orlando, Florida, USA) (ASSETS '10). ACM, New York, NY, USA, 279–280. <https://doi.org/10.1145/1878803.1878870>

## A DEMOGRAPHICS OF STUDY PARTICIPANTS

ID	Age/Gender	Visually impairment	Device
P1	24/F	Blind (total), acquired	smartphone
P2	24/F	Blind (some light perception), born	smartphone
P3	22/F	Blind (some light perception), born	smartphone, laptop
P4	25/M	Blind (total), born	smartphone
P5	24/M	Blind (some light perception), born	smartphone, laptop
P6	24/M	Low Vision (narrowed field), born	smartphone, laptop
P7	27/M	Low Vision (narrowed field), born	smartphone
P8	32/M	Blind (total), born	smartphone, laptop
P9	28/M	Blind (total), born	smartphone, laptop
P10	22/M	Blind (total), acquired	smartphone, laptop
P11	27/M	Blind (total), born	smartphone, laptop
P12	23/F	Low Vision (narrowed field), acquired	smartphone, laptop
P13	34/M	Blind (total), born	smartphone, laptop
P14	35/M	Blind (some light perception), acquired	laptop
P15	33/M	Blind (total), born	smartphone, laptop
P16	22/F	Blind (some light perception), born	smartphone, laptop
P17	27/M	Blind (some light perception), born	smartphone, laptop
P18	31/M	Blind (total), born	smartphone, laptop
P19	26/F	Blind (some light perception), born	smartphone, laptop
P20	44/F	Blind (total), acquired	smartphone, laptop

*P1 - P10: Chinese Users P11 - P20: North America Users*