Covering Covers

Characterization Of Visual Elements Regarding Sleeves

Yessin Beyhan

Supervisor: M.S. Pera

EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
January 24, 2023

Name of the student: Yessin Beyhan
Final project course: CSE3000 Research Project
Thesis committee: M.S. Pera, <Examiner>

Abstract

The cover of a children's book could play a crucial role in attracting the attention of young readers and influencing their decision to pick up and read the book. In this study, we aim to identify the factors that contribute to the appeal of children's book covers. Given that the target audience for children's books changes as children grow older, we divide our analysis into five age groups. We conduct a series of experiments in which we analyze three distinct aspects of covers: the visual attributes, visible objects, and implied story. Our results show that visual attributes play an important role in determining the appeal of a book. We also find that objects and themes present on a book cover vary between age groups. By understanding the elements that contribute to the appeal of children's book covers, we hope to improve recommendations for young readers and encourage a lifelong love of reading.

1 Introduction

Selecting the right book for a child can be a challenging task, especially given the vast number of options available [29], and the wide range of reading levels [15] and interests among young readers [2; 28]. There are many tools available to help adults choose books for themselves, such as book reviews and best seller lists, as well as recommendation systems (RS) based on content or collaborative filtering [32]. However, there is a lack of reliable RS for children's books [11]. This is particularly problematic given the importance of reading for children's cognitive development [3] and the role that books play in shaping their interests and values [31]. The scarcity of these RS can be attributed to the limited data available on children's user-system interactions, as well as a lack of research efforts focused on designing such systems [11].

Book covers play a crucial role in attracting the attention of young readers and influencing their decision to pick up and read the book [19; 34]. Previous research has shown that book covers can impact the book's appeal to children of different ages [16]. Some RS like Storytime [25] have focused on cover design as a key aspect, but little is known about the specific factors that contribute to the appeal of children's book covers and how these factors may vary as children grow.

Understanding these visual features may be crucial for creating effective and age-appropriate book recommendations, which is why we study the "*extent to which visual features of covers could contribute to the recommendation of books appealing to children of different ages*." For this, we examine the appeal of covers to children in five age groups: 0-3, 4-6, 7-9, 10-12, and 13+. These groups were chosen as they align with the reading developmental stages that children go through as they grow up [1; 13; 39; 47].

For analysis purposes, we turn to the Goodreads [43; 44], Amazon [18] and Open Library [21] datasets, which offer metadata and cover images on books that are age appropriate for children. It provides a diverse sample of children's books from different authors, publishers, and eras, ensuring that the results of the study can be applied to a wide range of children's book covers. To analyze the impact of book covers, we focus on three specific aspects: *visual attributes*, such as brightness and colorfulness, the *objects* visible on the cover, and the *story* that is implied in the cover. While there are studies that have looked at trends in visual attributes across age groups [24], this study is unique in its examination of objects and implied story in relation to age.

Overall, we found prominent patterns in the visual attributes, objects, and stories portrayed on children's book covers as they pertained to different age groups. We observed that older ages tend to have covers with darker colors and fewer objects depicted, while younger ages had covers with brighter colors and a greater variety of objects. Additionally, we found that the stories implied on the covers also change with age, with younger ages having covers that imply simpler and more fantastical stories, while older ages have covers that imply more complex and realistic stories. These findings suggest that there are distinct differences in the way book covers are designed to appeal to children of different ages.

By analyzing book covers that target specific age groups of children, we aim to gain insight into what drives the perceived appeal of children's literature. This in turn could be used to further the development of an effective RS for young readers, ultimately helping them discover the joy of reading.

2 Methodology

In this section, we describe in detail the methods, techniques, and experiments we conducted to attain our research goals. Further, we describe the age groups we use in our study; the data collection and preprocessing methods, as well as an explanation of the techniques used for analyzing the covers.

In the interest of transparency and reproducibility, the code for this project is publicly available on GitHub: https://github.com/Yessin111/Childrens-Book-Cover-Analysis. This repository contains all the scripts used to both process and analyze the data for this study. We encourage anyone who is interested in reproducing or building upon our work to use the provided code as a starting point.

2.1 Age groups

In order to investigate our research question, we cluster users into five different groups based on age. While it is possible to analyze every age separately, grouping them into broader categories allows for more meaningful comparisons, as similarities and differences are likely to be more pronounced than those between individual ages. Furthermore, a larger sample size per age group results in more precise measurements.

We define the age groups as: 0-3 years, 4-6 years, 7-9 years, 10-12 years, and 13 years and older. These age groups are chosen based on the stages of reading development, where the emphasis of reading education shifts from basic language skills for the youngest children to more complex concepts for older children. For books in the *0-3 age group* the focus is on building pre-reading skills, such as object recognition, language development and fostering a love of books and reading. They differ from books for older age groups as they are typically less text-heavy and rely more on visuals and physical elements to convey meaning [13]. Children at *ages 4-6* focus on developing basic language skills [39]. At *ages 7-9*, the emphasis shifts from learning how to read, to reading to learn. From this point forward, the ability to understand and interpret texts becomes a key focus in reading education [39]. Books intended for this age group have distinct characteristics in their organization, the way information is presented and the use of vocabulary [47]. As children advance to higher grades, the focus shifts to utilizing the information and knowledge gained from reading texts and books in other contexts [39]. Young adult literature is characterized by themes relevant to adolescents, such as relationships, identity, and finding one's place in the world. This genre is typically intended for readers aged *13-15 and older* [1].

2.2 Data description

The data for this study consists of children's book covers from various English publishing authors, spanning a range of reading levels from early childhood to young adult. We extract these books from the Goodreads dataset [43; 44], specifically from the "children" and "young adult" datasets. This results in a source, which we call *GRDS* (GoodReads DataSet), consisting of 217,480 instances, each including metadata like the title, author and country of publishing.

We extend the metadata from *GRDS* with the corresponding covers and ages. For this, we turn to the Amazon dataset [18], which contains cover images and the age range the book is appropriate for, allowing us to determine which age group(s) to assign each book to. Any books that fall into more than two age groups are deemed as having too broad of a range and are excluded from the study, as narrowing the focus to books that are more clearly targeted to a specific age group can provide more meaningful insights and conclusions about what makes book cover more appealing among that age group and why. To further refine *GRDS*, we also take into account the language of the book and exclude books that are not listed as being in English. By limiting the scope to English books, we can control for potential variations in cultural and societal factors that may influence the appeal of the covers across different languages. This measure reduces the instances in *GRDS* to 93,662. In order to expand the variety of book covers to analyze, we extend *GRDS* with the book cover images from Open Library [21]. By using cover images from multiple sources, we can reduce the potential influence of bias that might be present in just one source. An overview of *GRDS* is presented in Table 1. Because books can fall into more than one age group, the overall amount of covers for both sources is larger than the amount of unique books in *GRDS*.

Table 1: Distribution of books across age groups.

Age Group	Cover Source	
	Amazon	Open Library
0-3	21,351	16,830
4-6	44,966	36,556
7-9	43,575	35,762
10-12	30,801	24,909
13+	21,639	16,484
Overall	162,332	130,541

2.3 Experiment setup

We outline the setup for each of the experiments proposed to answer our research questions:

RQ1 How do the visual attributes of a cover influence the appeal a book has on children of different ages?

RQ2 How do the objects visible on a cover influence the appeal a book has on children of different ages?

RQ3 How does the implied story of a cover influence the appeal a book has on children of different ages?

Visual attributes

Drawing inspiration from Milton et al. [24], who also analyzed the visual attributes of covers, we consider the same lenses for this study (RQ1): dominant color, brightness, colorfulness, contrast, and entropy. To extract these traits, we use the Python Imaging Library [38] and OpenCV [7].

To determine the dominant color of each cover image, we use k-means clustering[1], as described in [37]. This is a simple, yet powerful algorithm that can identify patterns in large and complex datasets like images. By grouping the RGB values of all the pixels into nine clusters and finding the center of the largest cluster, we determine the most dominant color of the cover and match this color to a name using the CSS color module level 4 [42], which defines 139 different colors.

To measure the brightness of the cover images, we use a method inspired by the one described in [9], which calculates the brightness by converting the images to grayscale and taking the average pixel brightness. The colorfulness is calculated using the colorfulness metric methodology described in [12]. This metric defines colorfulness of an image as a series of calculations based on the mean and standard deviation of the RGB channels. For finding the contrast, we take the Root Mean Square contrast definition, inspired by the strategy presented in [30]. This method finds the contrast by calculating the standard deviation of the grayed image pixel intensities. The entropy is calculated using Shannon's entropy [36]. Entropy, in the context of cover images, is a measure of the disorder or randomness of the colors and other visual attributes of the image. We normalize all of these calculated visual attributes to a range of [0,1] to make it easier to understand and analyze the values.

Object detection

Objects in a cover can provide visual cues about the story and characters within the book and can influence a child's interest and perception of the book. Thus we probe objects automatically extracted from book covers (RQ2). However, the detection of objects is not a trivial task, as algorithms are typically trained on a small number of detection classes such as animals, vehicles, or people [22]. For example, an algorithm trained on animals is unable to detect furniture or tools.

As we aim to detect as many objects as possible on cover images, and traditional forms of object detection are limited to specific classes, we employ the use of zero-shot classification. This is a machine learning technique that allows for the identification of objects without the need for training data on those objects. Instead, the algorithm is trained on a set

[1]For the k-means, we set k=9, to account for the 7 colors of the rainbow plus black and white.

of source classes and then uses this training to classify new, unseen target classes. Specifically, we use zero-shot classification using 545 labels from Google's ImageNet dataset [17], which includes a wide range of object classes. By using this technique, we are able to accurately classify a variety of objects that may be present on the cover images, regardless of their presence in the training data. As the underlying architecture of our zero-shot classification, we use the OWL-ViT detection model [26], a vision transformer-based object detection model. We choose this model after manually testing many different models on a set of test cover images and finding that the OWL-ViT model gives the most accurate results[2].

The model generates a probability distribution for each of the 545 predefined labels, in a range of [0,1]. After manually looking at the same set of test images again, we set a threshold for certainty at 20%, meaning that we only consider the objects that the model is able to detect with a probability of 0.2 or higher. Any objects that are detected with a probability below this threshold are discarded. By setting this threshold, we are able to ensure that only the objects that the model is most confident in detecting are included in our analysis. This allows us to focus on the objects that are most likely to be present on cover images, minimizing the impact of false positives on our results.

Implied story

Lastly, we analyze the story conveyed by a book cover, which can give a more comprehensive understanding of how it appeals to children (RQ3). This analysis goes beyond just identifying objects, as it also examines characters and emotions.

To find the implied story of each cover image, we use the BLIP captioning model [20], which is trained on the COCO Captions dataset [8]. The COCO Captions set consists of over 330,000 images, each paired with at least 5 human-written captions describing the objects and scenes present in the image. The BLIP captioning model is a state-of-the-art machine learning model that generates descriptive captions for images.

One of the strengths of the BLIP model is its ability to understand the context between objects in an image. This allows it to generate more detailed and accurate captions, as it is able to describe not only the objects present in the image, but also their relationships and interactions with one another. This is important for our study, as understanding the implied story of the covers requires an understanding of the relationships between the visual elements and objects present on the covers, which is why we chose this model above others. In addition, the BLIP model is able to generate captions that are more human-like and easier to understand, as it takes into account grammatical structure and natural language patterns. This makes it an effective tool for generating captions that can be easily understood by humans and used for further analysis.

Much like we do for object detection, we manually inspect a sample of covers to ensure the accuracy of the results obtained from the model. We looked at the model's output to ensure that it was able to capture the main themes and ideas presented on the cover in a coherent manner.

[2] The test covers consist of 50 images, randomly chosen and equally distributed over all age groups from both cover sources.

2.4 Illustration

To illustrate the information we extract, we include an example cover in Figure 1 to demonstrate the results of using the methods described in Section 2.3.

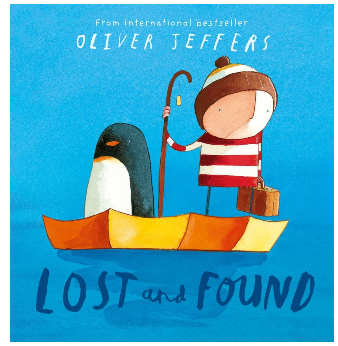

Figure 1: Cover of "Lost and Found" by Oliver Jeffers.

Visual attributes

The values for the extracted visual attributes are:

- Dominant Color: (13, 172, 226)
- Dominant Color Name: deepskyblue
- Brightness: 0.615
- Colorfulness: 0.519
- Contrast: 0.171
- Entropy: 0.955

The dominant color –deepskyblue– is fitting as the sea is a large part of the cover. The brightness and colorfulness values reflect the presence of both darker and lighter colors in the image, such as the sea and the depicted characters respectively. The contrast value is relatively low, which is appropriate given that the image features a lot of the same color. The illustrator made use of a lot of static, meaning small variations in color, which is reflected in the high entropy value.

Object detection

The objects extracted using the object detection are:

- Poster: 0.617
- Penguin: 0.485
- Clothing: 0.442
- Person: 0.428
- Hat: 0.274
- Paddle: 0.267
- Canoe: 0.24
- Mammal: 0.202

The penguin, which is a key feature of the cover, is detected with a high confidence score. Additionally, the person is also detected with a high confidence score. It is also notable that the upside down umbrella, functioning as a boat, got split into both a paddle and a canoe, which is an acceptable result for the purposes of object analysis.

Implied story

The result for extracted implied story is: "A picture of a penguin and a man in a boat". The caption identifies the key elements present in the image. It is worth noting that, as mentioned above, the upside down umbrella functioning as a boat was correctly identified as one, which is beneficial for analysis purposes as it allows for a more clear understanding of the objects present in the image and their relationship to each other. Overall, this is an accurate caption that effectively captures the essence of the image and its contents.

2.5 Analysis

Ultimately, we examine the trends emerging from the obtained data. We compute various statistics across different age groups to understand how visual features of covers contribute to the recommendation of books appealing to children of different ages. We also compute statistical significance tests to compare and contrast the trends across different age groups and sources.

Visual attributes

For the visual attributes, we examine the trends of dominant color usage across different age groups and sources. Specifically, we compute the proportion of all dominant colors present per age group. We then analyze the patterns emerging from these percentages. For the remaining attributes we perform a series of statistical tests to determine if there are any significant differences in their values across the different age groups and sources. We use a one-way ANOVA (from the SciPy [41] library) to determine whether there is a significant difference in the means of the attribute values for each age group. We then use the Turkey HSD test (from the statsmodels [35] library) to identify which specific pairs of age groups show significant differences in the attribute values.

Object detection

For the Object Detection analysis, we perform three types of analysis to understand the prevalence and relationship between objects found on children's book covers across different age groups: frequency analysis, correlation analysis, and co-occurence analysis.

Frequency analysis. We analyze the frequency by counting the number of times each object appears across all covers and age groups and comparing the frequency of objects across the different age groups, which shows us which objects are more prominent in each age group. For analysis we take a qualitative approach and examine patterns in the data.

Correlation analysis. Correlation analysis shows which objects are *relatively* more prominent in each group. We perform this analysis by treating objects as predictors and using their frequency to establish a pattern between the objects and the age groups. We use Pearson correlation (from the Pandas [23] library), as it is a widely used measure of the linear association between two variables. This enables us to identify which objects are indicative of each age group.

Co-occurence analysis. We analyze co-occurrence by identifying which objects frequently appear together on covers within the same age group. For this, we use point-wise mutual information (PMI) [6], often used in natural language processing to measure the association between words. The statistical association is measured between two variables, and is calculated as the logarithm of the ratio of the probability of the two elements co-occurring to the probability of each element occurring individually.

Implied story

To analyze the content of the captions in the children's books, we use Named Entity Recognition (NER) and Topic Modeling. To perform these analyses we make use of the natural language processing libraries spaCy [14] and Gensim [33] respectively.

Named Entity Recognition. NER is a technique that helps to extract and classify named entities in a text. It works by analyzing the structure and context of words within a sentence and comparing them to a pre-defined set of categories, such as people, organizations, and locations. We use the spaCy library, which uses a combination of rule-based and statistical approaches to identify named entities in the captions. It first tokenizes the text, breaking it down into individual words and their associated grammatical information. Then, it uses this information to identify and classify named entities within the text. The library is able to identify the following labels:

- `CARDINAL`: Numeric values not under another label.
- `DATE`: Absolute or relative dates or periods.
- `EVENT`: Named hurricanes, battles, wars, sports events, etc.
- `FAC`: Buildings, airports, highways, bridges, etc.
- `GPE`: Countries, cities, states.
- `LANGUAGE`: Any named language.
- `LAW`: Named documents made into laws.
- `LOC`: Non-GPE locations, mountain ranges, bodies of water.
- `MONEY`: Monetary values, including unit.
- `NORP`: Nationalities or religious or political groups.
- `ORDINAL`: "first", "second", etc.
- `ORG`: Companies, agencies, institutions, etc.
- `PERCENT`: Percentage values.
- `PERSON`: People, including fictional.
- `PRODUCT`: Objects, vehicles, foods, etc. (Not services.)
- `QUANTITY`: Measurements, as of weight or distance.
- `TIME`: Times smaller than a day.
- `WORK_OF_ART`: Titles of books, songs, etc.

We add the `ANIMAL` label to our NER model by checking all the captions against a list of 221 animal names that we extract from Wikipedia [46].

To determine differences in the distribution of the identified labels in each of the age groups, we conduct a chi-squared test using the SciPy library. This enables us to determine if certain labels were more or less prominent in certain age groups and to identify patterns in the captions of books targeted towards different age groups.

Topic modeling

Topic modeling works by identifying patterns in the co-occurrence of words in a set of texts. It uses statistical techniques to identify which words are likely to occur together in the same context, and groups them into topics. The algorithm starts by representing each text as a high-dimensional vector of word counts, and then uses Latent Dirichlet Allocation (LDA) [5] to identify the underlying topics. LDA is a generative probabilistic model that assumes that each document is a mixture of a small number of topics, and each topic is a probability distribution over the words in the vocabulary.

Using topic modeling, we identify the top 10 themes in the captions, each represented by the 10 most relevant words, accounting for stopwords using Natural Language Toolkit library [4]. We then use this information to categorize every caption into its most fitting topic. By focusing on the 10 most relevant words, we are able to capture the main ideas and concepts associated with each topic while avoiding including less important words that would add noise to the analysis. It is important to note that Topic modeling is an unsupervised method, meaning that it does not use predefined labels. Instead, it uses the patterns of words within the text to identify topics, meaning the results are not always easy to interpret.

We examine the words associated with each topic and find a general theme associated with the 10 most relevant words. A general theme can be found by examining these words as a whole and identifying patterns or trends in the meanings or associations. For example, if the words for a topic are almost all related to animals, it can be inferred that the topic's theme is related to animals. We investigate how the topic distribution changes as age increases, to understand how themes evolve as the intended audience changes. We conduct a chi-squared test to determine whether the shifts in topics are statistically significantly different across different age groups.

3 Results

Here, we discuss the outcomes emerging from the experiments conducted to answer our RQs.

3.1 Visual attributes

We analyze the visual attributes to see whether there are significant differences between the different ages. We discuss salient findings below. Detailed outcomes of the statistical tests are included in Appendix B, Tables 3, 4, 5 and 6.

Dominant color. We first consider the distribution of dominant colors for the cover images of both Amazon and Open Library. It should be noted that for the purpose of clearly displaying and to obtain more meaningful results, the 139 color names have been grouped into 9 clusters, one for each color of the rainbow plus black and white.

As shown in Figure 2, white and black are the most dominant colors for both sources, across all age groups. Combined, they make up 50% to 70% of the dominant colors in all of the age groups. This suggests that a significant proportion of children's book covers are designed with a neutral color palette, which could be a deliberate choice to appeal to a wide range of age groups.

When looking at the distribution of colors by age groups, it is clear that, for both sources, the percentages of white colors decrease and the percentages of other colors tend to increase as age increases. This indicates that as children grow older, they may be more likely to be drawn to book covers with more varied and dynamic color palettes. It is worth noting that for each age group for both sources the overall shift was statistically significant for every dominant color.

It is also interesting to note there are clear patterns in the change of color percentages as age increases, across both Amazon and Open Library. On both platforms, the percentages of orange, yellow and blue tend to decrease as age increases, while the percentages of green and violet tend to increase. When it comes to the fluctuating colors, the percentages of red and indigo show an upward trend until age groups 7-9 and 10-12 on both sources, after which it decreases again. This suggests that both red and indigo colors may be particularly appealing to children in the middle age groups.

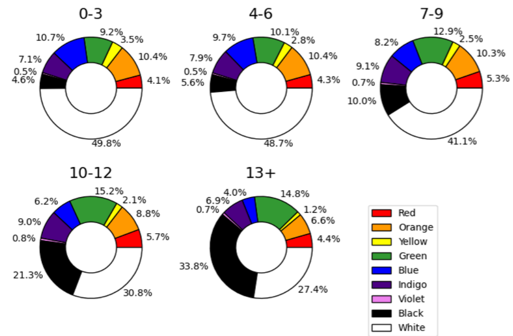

(a) Amazon

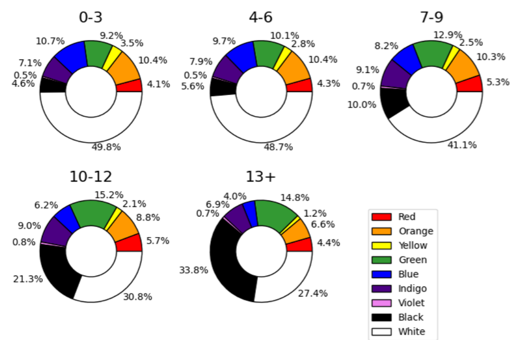

(b) Open Library

Figure 2: Dominant colors of cover images per age group.

Brightness. The brightness of the cover images is shown in Figure 3. We see that there are significant differences in brightness between all pairs of age groups across both sources. The tests indicate that every group is overall darker than its predecessors. The implications of this finding suggest that as children grow older, the book covers that are marketed to them become darker in color. This may be because publishers and authors believe that older children are ready for more mature and serious themes, which are often conveyed through darker colors. Additionally, darker colors may be used to create a sense of mystery or intrigue, which can be appealing to older children.

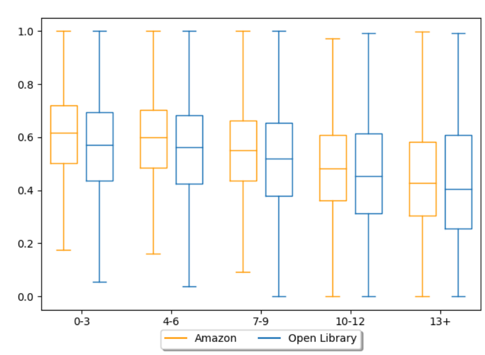

Figure 3: Brightness of covers per age group.

Colorfulness

The colorfulness of the cover images is shown in Figure 4. There are significant differences in the colorfulness of book covers across the different age groups: the colorfulness of book covers decreases with every age group, implying that book covers aimed at younger children tend to be more vibrant and colorful in comparison to those aimed at older children. Additionally, it could also suggest that as children age, they are more drawn to book covers that are more subdued and less bright, which could reflect a preference for more mature and realistic themes.

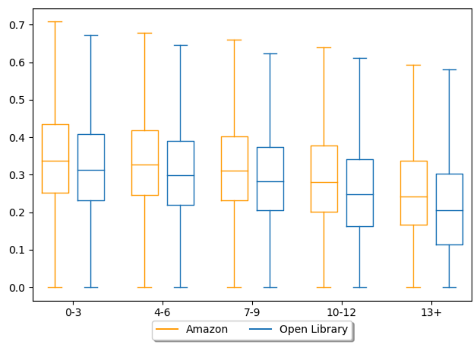

Figure 4: Colorfulness of covers per age group.

Contrast

The contrast of the cover images is shown in Figure 5. The contrast attribute also shows significant differences between all pairs of groups, except for ages 0-3 and ages 4-6 for both Amazon and Open Library, and ages 0-3 and 13+, 4-6 and 13+, and 7-9 and 10-12 for Open Library. Overall, the results of the contrast attribute suggest that there may be a trend towards increasing contrast in book covers as the age group increases, with the exception of few age groups. This could suggest that publishers might be using higher contrast in covers for older children as it could be more attention-grabbing and appealing to them. However, further research is needed to confirm this trend and to determine the precise nature of the relationship between contrast and age group.

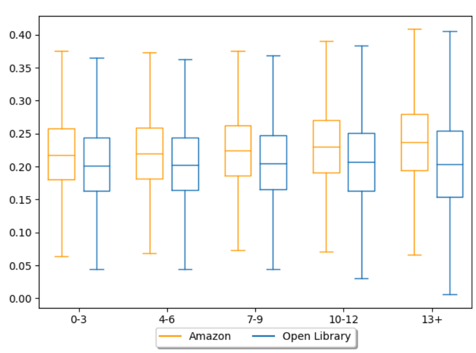

Figure 5: Contrast of covers per age group.

Entropy

The entropy of the cover images is shown in Figure 6. We see a statistically significant difference in the entropy values among the different age groups in both datasets. For both Amazon and Open Library, there is a significant increase in entropy from ages 0-3 to ages 7-9, after which there is a significant decrease from ages 7-9 to ages 13+. This is the only attribute where we observed both statistically significant increases and decreases, suggesting that there is some complex relationship between age and entropy in book cover images.

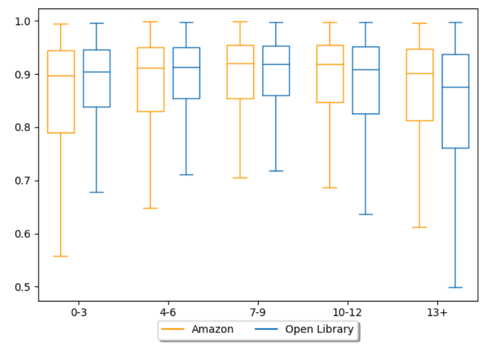

Figure 6: Entropy of covers per age group.

3.2 Object detection

We present the results of our frequency and correlation analysis of the objects detected below. Due to the large number of labels used for the zero-shot detection, we focus on discussing the most notable results. Details of statistical tests and results can be found in Appendix C, Tables 7, 8 and 9.

Frequency

For each age group and both sources, the most prominent object is "poster," the second one is "clothing," and the third and fourth objects are "person" and "face," in either order. The reason for the high frequency of the "poster" object - 87% - is likely due to the object detection being trained on real-life photographs, and interpreting the book covers as posters. The prevalence of the "clothing", "person", and "face" objects, which are all associated with people, indicates that human figures are a common element on covers of children's books.

In the 0-3 and 4-6 age groups, nature is commonly represented with objects such as "animal", "chicken", "rabbit", "tree", and "flower". Additionally, toys also make a minor appearance, with objects like "toy" and "teddy bear". However, as the age group increases, we see a shift towards people-related objects such as "woman", "man", "girl", and personal accessories. Objects like "footwear" and "hat", which were also present in earlier age groups, become more common. This trend continues in the 10-12 and 13+ age groups, where we see an even stronger focus on people and their personal features and belongings, like "lipstick", "hair", and "clothing". This suggests that children's books aimed at younger readers may be more likely to feature illustrations of animals and nature, while books aimed at older readers may be more likely to feature illustrations of people.

Correlation
Correlation analysis gives an insight into which objects are *relatively* more prominent in each group. When the correlation coefficient is close to 1, it means that there is a strong positive correlation between the two variables; when one variable increases, so does the other. Conversely, there is a strong negative correlation when the coefficient is close to -1 . In our case, the correlation values are all close to 0, ranging between -0.21 and 0.34, meaning that there is no strong correlation between the frequency of the objects and the age groups. In other words: the objects on covers can not be used as good predictors of which age group a book is targeted at.

Despite the weak correlations, the objects that are present in the correlation analysis do align with the patterns observed in the frequency analysis. We see a positive correlation between objects related to animals and toys for the youngest age group and a positive correlation between objects related to people for the oldest age groups. Additionally, the negative correlation for both ends of the age groups is inverse to each other, meaning that there is a negative correlation between objects related to people for the youngest age group, and a negative correlation between objects related to animals and toys for the oldest age group.

Co-occurence
The results, although showing some trends in the types of objects that are commonly found together on book covers, do not provide any clear information or insights into how these object pairs may be related to the appeal of a book to different age groups. While there is a shift from objects concerning toilet items such as "bidet" and "toiler paper" in the younger age groups to more food and kitchen-related items such as "oyster" and "sushi" in the older age groups, these results are neither in line with the findings from the frequency and correlation analysis, nor informative enough to draw any meaningful conclusion from.

3.3 Implied story

NER and topic modeling analysis on the captions of the children's books are outlined below. We only show and discuss salient findings. Detailed outcomes of the analysis are included in Appendix D, Tables 12, 13 and 14.

Named Entity Recognition
The results of the NER analysis show that, for both sources, the number of mentions of `ANIMAL`, `CARDINAL`, `DATE`, `GPE`, `NORP`, `ORDINAL`, `ORG`, `PERSON` and `TIME` all have significant differences in their frequency of mention across age groups. `FAC`, `PRODUCT` and `QUANTITY` are all significant for Amazon but not for Open Library. On the other hand, the number of mentions of `EVENT`, `LANGUAGE`, `LAW`, `LOC`, `MONEY`, `PERCENT` and `WORK_OF_ART` did not show any significant changes with age. Overall, almost all of the labels were mentioned less as age increased. One possibility for this is that the books with older age groups may be more focused on serious topics and therefore less likely to mention entities such as animals, locations, and events, which are more commonly associated with lighter, more whimsical content. This suggests that as books become targeted towards older age groups, the covers tend to be less visually rich. It is worth noting that all of these trends were consistent across both the Amazon and Open Library datasets, which would indicate that they may be generalizable to other book cover captions as well.

Topic Modeling
We analyze the top 10 words for each of the 10 distinct topics. The top words associated with each topic give some indication of the theme or subject matter of each topic:

- `Topic 1`: Children and their activities, with words like "children", "playing", and "summer" indicating a focus on childhood and social interaction.
- `Topic 2`: Stories and narratives, with words like "words", "story", and "secret" indicating a focus on written content and possibly fiction.
- `Topic 3`: Nature, with words like "bear", "snow", "flower" and "life" indicating a focus on flora and fauna.
- `Topic 4`: Art and creation, with words like "painting", "drawing", "bird" and "child" indicating a focus on creative expression and visual representation.
- `Topic 5`: Adventure, with words like "horse", "night", and "adventures" indicating a focus on exciting experiences.
- `Topic 6`: Fantasy and imagination, with words like "flying", "sky", "world" and "princess" indicating a focus on magical and imaginary elements.
- `Topic 7`: Color, with words like "black", "white", "red" and "blue" indicating a focus on the visual aspect of books.
- `Topic 8`: Animals, with words like "dog", "cat" and "dragon" indicating a focus on pets and wild animals, both real and imaginary. This topic can also be associated with young femininity, with words like "girl," "pink," "little," and "dress" indicating a focus on young girls.
- `Topic 9`: Relaxation and romance, with words like "beach", "couple" and "heart" indicating a focus on leisure and love.
- `Topic 10`: School and youth, with words like "young", "boy", "school" and "field" indicating a focus on education.

The variation in the distribution of captions across the 10 different topics as identified by the topic modeling algorithm is significant across the different age groups. For example, `Topic 1` is more prevalent in the age groups 0-3 and 4-6, with a decrease in frequency as age increases. This finding is in line with the set of words identified before, which suggested that this topic is related to themes of childhood and social interaction. Similarly, `Topic 8` is also more prevalent in the age groups 0-3 and 4-6, with a decrease in frequency as age increases. This is in line with the themes associated with this topic: animals and young femininity.

`Topic 9` is more prevalent in 10-12 and 13+ age groups, with a decrease in frequency as age decreases. This aligns with our expectations, as the topic is related to themes of relaxation and romance. `Topic 10`, associated with education, is more prevalent in the 7-9, 10-12 and 13+ age groups, with a decrease in frequency as age decreases. This finding is in line with our expectation, as children in this age group are able to relate to the themes of growing up, school, and education.

Overall, the distribution of captions across topics changes with age, which can provide insight into how the themes and content of the covers change with the intended audience.

4 Discussion

For visual attributes, the brightness, colorfulness, entropy, and for the most part contrast too, varied significantly across different age groups, with covers targeted to younger age groups tending to be brighter and more colorful. Additionally, as the target audience of the books shifts to older age groups, the dominant colors on the covers change, with an increase in green and violet and a decrease in orange, yellow and blue. As previously mentioned, there are other studies that have also looked at trends in visual attributes across age groups. Our findings are in line with theirs with results that are more statistically significant, highlighting the importance of visual attributes in book cover design and the impact it has on the appeal of a book to children of different ages.

Our object detection analysis reveals that certain objects are more prevalent in certain age groups. For example, the frequency analysis shows that objects related to animals and nature are more common in covers targeted towards younger age groups, while those related to people and accessories are more common in covers targeted towards older age groups. However, while the correlation analysis supports this conclusion, the correlation values are not significant enough to use these objects as reliable indicators for the age groups. Unfortunately, the co-occurence data does not provide much insight either. While we see a gradual shift from pairs concerning toilet items to more kitchen-related items, this data seems to be of limited use in determining the appeal of a book based on the objects present on its cover. All things considered, these results suggest that even though the objects present on a book cover vary between the different age groups, they can not be used to reliably indicate the intended age group of the book.

The results of the implied story research show shifts in themes across the age groups. The NER reveals that labels like `ANIMAL`, `CARDINAL` and `GPE` are mentioned more in covers targeted towards younger age groups, with almost none of the labels being mentioned more in covers targeted towards older age groups. Additionally, topic modeling shows that covers portraying topics related to "nature" and "adventure" are more common in covers targeted towards younger age groups, while topics related to "school" and "education" are more common in covers targeted towards older age groups. Overall, these outcomes show that the implied story on a cover plays a role in appealing to children of different ages.

The results of this research have noteworthy implications for the development of comprehensive and reliable RS for children's books. It highlights the importance of considering the visual attributes, visible objects, and implied story of a book cover, which can be used to develop more effective and age-appropriate book recommendations. For example, by understanding that covers targeted towards younger age groups tend to be brighter and more colorful, a RS could take this into account and prioritize covers that match these attributes for younger readers. Similarly, by understanding that certain objects are more appealing to certain age groups, a RS could take this into consideration and prioritize books with those objects for the corresponding age groups. All in all, the results contribute to the understanding of what makes a book cover appealing to children of different ages and can be used to improve the accuracy and effectiveness of book recommendations for children.

Traditional RS leverage readability levels, topics, or ratings when recommending books, which are based on content or historical user-system interactions. However, the findings from this work offer a different perspective that could complement existing solutions. By incorporating the visual attributes, visible objects, and implied story of a book cover, RS can provide a more personalized experience for the user while balancing the trade-off between personalization and privacy. These findings can also help address the cold start problem, the challenge RS face to provide accurate recommendations for new users or new items, by providing perspective on recommending books based on visual attributes.

5 Conclusions, Limitations and Future Work

With this work, we aimed to understand the appeal of book covers on children of different ages. We focused on three aspects of covers: their visual attributes, the objects depicted, and the implied story. By analyzing a large dataset of covers, we identified trends and patterns that may be of interest to researchers and practitioners in the field of children's literature.

One limitation of this research are the age group boundaries that were used. While we attempted to divide the ages of the books into five distinct groups, there may still be some overlap in the reading preferences of children within these groups. In the future, we would like to explore more fine-grained age groups or different ways of grouping the books. Another limitation is that we used zero-shot object detection for detecting objects on covers. While using zero-shot object detection allows us to identify a wide range of objects on cover images, it is not as accurate as using a regular form of object detection that is trained specifically on objects often present in children's book covers. In the future, we would consider collecting and labelling a large dataset of children's book covers to train a more specialized object detection algorithm to improve the accuracy results. Lastly, this research was based on external sources for the cover images, which could introduce bias or unreliability to the data. In order to ensure the validity of the data, multiple sources were used, however, there is still a possibility that the results are not representative of the entire population of children's books.

Overall, our study provides insights into the factors that influence the appeal of children's book covers to different age groups. However, there are many areas for future research to explore. Further studies could examine the impact of different design elements on children's reading experiences by incorporating additional features of book covers (e.g. font [27] or layout) into the analysis to better understand how these elements contribute to the appeal of a book cover. Additionally, this work could also serve as a template for cover exploration for other user groups like adults [45] or even be applied to completely different areas of recommendation, such as analyzing the appeal of movie [10] or music album covers [40]. We hope that this research will serve as a starting point for these further investigations and ultimately lead to the development of more effective and personalized recommendation systems.

6 Responsible Research

In conducting this research, we took several steps to ensure that it was conducted in a responsible and ethical manner. We prioritize transparency in the data acquisition and analysis process. The data used in this study was obtained from publicly available sources. Additionally, we made our code publicly available to ensure reproducibility of our results. This allows other researchers to verify our methods and replicate our experiments if desired. Additionaly, we properly cited all sources and give credit to the appropriate parties. We made sure to properly attribute the data we used, including the datasets from Goodreads and Open Library, as well as any third-party tools or resources that were used in the course of our research.

Overall, we ensured that this research was conducted in a responsible and ethical manner in our effort to make a valuable contribution to the fields of both children's literature and recommender systems.

7 Acknowledgements

I would like to express my gratitude to my proofreaders for their valuable notes and feedback. I would also like to extend a special thank you to my supervisor and professor, M.S. Pera, for her guidance and support throughout this project. Her insight and expertise have been invaluable. I am also grateful to my girlfriend for her input and for her support during the time I devoted to this project. Their contributions have greatly improved the quality of this research.

References

[1] Linda Ackermans. Een decennium young adult literatuur in nederland (2009-2019): een stand van zaken. *Tijdschrift voor Nederlandse Taal-en Letterkunde*, 137(1):45–72, 2021.

[2] Richard E Ahl and Yarrow Dunham. Have your cake, and your asparagus, too: Young children expect variety-seeking behavior from agents with diverse desires. *Cognitive Development*, 54:100882, 2020.

[3] Richard C Anderson, Larry L Shirey, Paul T Wilson, and Linda G Fielding. Interestingness of children's reading material. In *Aptitude, learning, and instruction*, pages 287–300. Routledge, 2021.

[4] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009.

[5] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.

[6] Gerlof Bouma. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL*, 30:31–40, 2009.

[7] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

[8] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.

[9] cmcginty. What are some methods to analyze image brightness using python?, Aug 2010.

[10] Cecilia Fagerholm. The use of colour in movie poster design: An analysis of four genres. *N.A.*, 2009.

[11] Jerry Alan Fails, Maria Soledad Pera, Franca Garzotto, and Mirko Gelsomini. Kidrec: Children & recommender systems: Workshop co-located with acm conference on recommender systems (recsys 2017). In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pages 376–377, 2017.

[12] David Hasler and Sabine E Suesstrunk. Measuring colorfulness in natural images. In *Human vision and electronic imaging VIII*, volume 5007, pages 87–95. SPIE, 2003.

[13] Elizabeth A Hasson. "reading" with infants and toddlers. *Day Care and Early Education*, 19(1):35–37, 1991.

[14] Matthew Honnibal and Ines Montani. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear, 2017.

[15] Charles Hulme and Margaret J Snowling. *Developmental disorders of language learning and cognition*. Blackwell History of Literature. Wiley-Blackwell, Chichester, England, December 2008.

[16] Leigh Ann Jones. *Do "Quick Picks" book covers appeal to middle school readers? The impact of covers on self-selection of fiction books in the middle school library setting*. Texas Woman's University, 2006.

[17] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981, 2020.

[18] Himabindu Lakkaraju, Julian McAuley, and Jure Leskovec. What's in a name? understanding the interplay between titles, content, and communities in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7, pages 311–320, 2013.

[19] Monica Landoni and Elisa Rubegni. Aesthetic relevance when selecting multimedia stories. *International Workshop on Children & Recommender Systems*, 2017.

[20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022.

[21] Open Library. Open library. https://openlibrary.org/, 2015.

[22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[23] Wes McKinney et al. Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference*, volume 445, pages 51–56. Austin, TX, 2010.

[24] Ashlee Milton, Levesson Batista, Garrett Allen, Siqi Gao, Yiu-Kai D Ng, and Maria Soledad Pera. "don't judge a book by its cover": Exploring book traits children favor. In *Fourteenth ACM Conference on Recommender Systems*, pages 669–674, 2020.

[25] Ashlee Milton, Michael Green, Adam Keener, Joshua Ames, Michael D Ekstrand, and Maria Soledad Pera. Storytime: eliciting preferences from children for book recommendations. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 544–545, 2019.

[26] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with vision transformers. *arXiv preprint arXiv:2205.06230*, 2022.

[27] Taiga Miyazono, Brian Kenji Iwana, Daichi Haraguchi, and Seiichi Uchida. Font style that fits an image–font generation based on image context. In *International Conference on Document Analysis and Recognition*, pages 569–584. Springer, 2021.

[28] Kathleen AJ Mohr. Children's choices for recreational reading: A three-part investigation of selection preferences, rationales, and processes. *Journal of Literacy Research*, 38(1):81–104, 2006.

[29] N.A. *LIBRARY AND BOOK TRADE ALMANAC 2022: Formerly the Bowker annual*. N.A., 2022.

[30] Franco Piccolo. How to calculate the contrast of an image?, Nov 2019.

[31] Husni Rahim and Maila Dinia Husni Rahiem. The use of stories as moral education for young children. *International Journal of Social Science and Humanity*, 2(6):454, 2012.

[32] Sushama Rajpurkar, Darshana Bhatt, Pooja Malhotra, MSS Rajpurkar, and MDR Bhatt. Book recommendation system. *IJIRST–International Journal for Innovative Research in Science & Technology*, 1(11):314–316, 2015.

[33] Radim Rehurek and Petr Sojka. Gensim–python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, 3(2), 2011.

[34] Kara Reuter. Assessing aesthetic relevance: Children's book selection in a digital library. *Journal of the American Society for Information Science and Technology*, 58(12):1745–1763, 2007.

[35] Skipper Seabold and Josef Perktold. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*, 2010.

[36] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

[37] Shivam Thakkar. Dominant colors in an image using k-means clustering, Jul 2019.

[38] P Umesh. Image processing in python. *CSI Communications*, 23, 2012.

[39] Paul Van den Broek, Jolien M Mouw, and Astrid Kraal. Individual differences in reading comprehension. In *Handbook of individual differences in reading*, pages 156–168. Routledge, 2015.

[40] Tara Venkatesan, Qian Janice Wang, and Charles Spence. Does the typeface on album cover influence expectations and perception of music? *Psychology of Aesthetics, Creativity, and the Arts*, 2020.

[41] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[42] W3C. Css color module level 4. https://www.w3.org/TR/css-color-4/#named-colors, Nov 2022.

[43] Mengting Wan and Julian McAuley. Item recommendation on monotonic behavior chains. In *Proceedings of the 12th ACM conference on recommender systems*, pages 86–94, 2018.

[44] Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian McAuley. Fine-grained spoiler detection from large-scale review corpora. *arXiv preprint arXiv:1905.13416*, 2019.

[45] Alexis Weedon. In real life: book covers in the internet bookstore. In *Judging a Book by Its Cover*, pages 139–148. Routledge, 2016.

[46] Wikipedia. List of animal names - Wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=List_of_animal_names, 2023.

[47] Michael BW Wolfe. Memory for narrative and expository text: independent influences of semantic associations and text organization. *Journal of Experimental Psychology: Learning, memory, and cognition*, 31(2):359, 2005.

A Dominant colors

Table 2: Dominant colors in cover images per age group.

	Amazon					Open Library				
Color	0-3	4-6	7-9	10-12	13+	0-3	4-6	7-9	10-12	13+
aliceblue	88	234	184	82	60	44	124	105	46	25
antiquewhite	133	323	307	188	129	102	257	384	375	303
aquamarine	25	35	21	16	10	15	27	15	12	7
azure	3	8	10	7	3	0	6	5	5	2
beige	195	439	406	204	69	149	323	289	160	64
bisque	112	253	212	105	54	58	132	122	67	38
black	949	2,457	4,275	6,504	7,290	1,137	2,958	4,826	6,278	5,997
blanchedalmond	85	160	127	67	20	45	92	77	36	16
blue	2	5	7	4	0	1	6	8	4	1
blueviolet	0	4	7	4	1	1	2	4	2	0
brown	84	214	238	177	118	69	180	201	133	84
burlywood	174	447	438	235	104	123	308	297	174	67
cadetblue	131	291	276	158	92	130	257	218	115	59
chartreuse	2	4	3	0	0	3	2	2	0	0
chocolate	166	327	297	212	121	123	247	231	138	71
coral	77	148	121	75	39	46	81	67	44	22
cornflowerblue	261	505	432	211	92	145	291	236	103	48
cornsilk	73	129	94	53	24	44	70	49	26	15
crimson	255	454	464	344	194	132	253	221	147	87
cyan	2	4	4	3	3	1	1	4	5	1
darkblue	5	14	11	7	1	4	14	13	7	3
darkcyan	214	360	268	170	81	133	235	180	92	44
darkgoldenrod	15	38	47	24	7	113	218	189	92	40
darkgray	301	684	692	458	255	573	1,220	1,106	634	390
darkgreen	21	57	73	56	24	34	86	94	48	7
darkkhaki	276	632	570	282	134	232	546	458	184	67
darkmagenta	2	7	12	13	4	1	4	8	9	4
darkolivegreen	204	661	844	537	254	215	595	709	443	167
darkorange	33	62	74	53	27	36	61	48	27	18
darkorchid	4	4	6	4	1	1	1	5	4	1
darkred	11	32	42	37	28	56	105	119	83	52
darksalmon	48	88	95	62	33	31	64	64	38	20
darkseagreen	173	340	267	116	35	115	248	211	101	35
darkslateblue	385	807	782	492	233	313	633	576	329	132
darkslategray	960	2,454	3,554	3,493	2,651	724	1,919	2,408	2,070	1,447
darkturquoise	71	132	110	82	47	23	46	45	29	13
darkviolet	1	1	2	1	0	0	1	2	1	0
deeppink	7	18	29	43	50	4	10	16	23	27
deepskyblue	81	134	113	78	37	44	71	66	38	16
dimgray	271	728	929	709	388	258	662	777	536	263
dodgerblue	98	163	160	101	27	53	100	90	51	16
firebrick	140	286	315	219	119	161	325	338	203	95
floralwhite	178	395	318	138	84	102	241	239	161	149
forestgreen	53	94	79	42	11	47	79	69	32	6
gainsboro	503	1,145	1,033	626	438	573	1,336	1,174	724	519
ghostwhite	169	327	238	124	117	74	148	123	84	61
gold	397	613	515	303	118	224	334	247	140	73
goldenrod	217	416	348	190	87	181	328	284	155	52
gray	194	519	578	343	186	355	825	796	491	286
green	8	12	9	3	0	9	16	11	3	2
greenyellow	34	58	50	35	18	23	40	32	20	8
honeydew	40	99	84	34	11	33	83	70	28	9
hotpink	23	46	55	49	31	20	32	29	29	24
indianred	117	256	264	166	88	77	191	213	128	56
indigo	8	19	37	42	18	12	27	39	36	15
ivory	91	144	98	51	22	23	51	45	28	15
khaki	739	1,561	1,241	519	181	457	997	779	317	105
lavender	477	1,009	772	394	240	261	583	430	209	132

lavenderblush	18	59	53	34	27	12	37	31	17	14
lawngreen	1	8	7	5	3	2	3	3	1	1
lemonchiffon	93	171	132	55	15	71	142	103	37	9
lightblue	303	582	436	204	77	155	338	264	114	43
lightcoral	22	39	52	35	14	12	21	25	15	9
lightcyan	46	111	74	38	25	35	88	53	16	11
lightgoldenrodyellow	80	161	118	49	20	56	116	86	38	9
lightgray	262	575	558	398	329	393	853	827	579	368
lightgreen	15	28	28	15	5	12	23	16	9	5
lightpink	69	134	137	98	66	48	87	82	48	26
lightsalmon	24	52	50	30	23	15	32	25	9	4
lightseagreen	144	292	275	168	77	93	180	152	74	37
lightskyblue	78	137	105	45	16	40	73	54	26	8
lightslategray	169	378	395	244	132	196	407	364	223	99
lightsteelblue	277	646	582	336	149	202	504	428	199	82
lightyellow	33	68	53	27	9	18	24	24	19	5
lime	4	7	4	1	0	2	3	2	1	0
limegreen	23	39	27	11	6	14	25	21	11	3
linen	268	636	579	306	216	288	724	663	485	404
magenta	2	4	4	3	2	2	3	4	4	2
maroon	66	150	271	264	152	95	223	279	223	128
mediumaquamarine	104	196	168	96	56	67	147	124	51	23
mediumblue	6	11	8	7	4	4	12	8	7	4
mediumorchid	6	12	14	10	3	5	11	6	3	4
mediumpurple	34	72	65	40	17	24	54	42	18	11
mediumseagreen	70	128	97	46	16	45	90	64	19	6
mediumslateblue	4	7	7	6	0	1	2	2	2	0
mediumspringgreen	0	0	1	1	0	0	1	2	3	0
mediumturquoise	176	305	220	128	66	100	160	124	81	34
mediumvioletred	22	61	89	67	45	15	38	61	56	36
midnightblue	222	507	615	532	349	260	570	588	393	214
mintcream	135	327	247	93	48	108	218	165	76	39
mistyrose	59	112	86	54	39	39	72	50	35	27
moccasin	108	232	218	112	34	47	117	109	47	6
navajowhite	71	137	116	57	26	30	68	70	30	12
navy	3	9	11	11	4	5	8	9	8	2
oldlace	154	348	280	115	54	88	199	192	117	81
olive	1	7	15	8	2	39	89	78	41	21
olivedrab	113	251	221	90	30	98	227	207	95	28
orange	52	105	87	49	18	42	73	57	29	14
orangered	52	106	101	66	36	48	75	70	51	20
orchid	2	10	12	7	7	2	10	10	9	6
palegoldenrod	242	544	460	195	64	158	388	344	144	36
palegreen	6	16	14	9	3	2	11	16	10	1
paleturquoise	125	255	187	92	46	60	144	113	45	18
palevioletred	74	135	134	104	70	45	88	108	76	39
papayawhip	28	53	38	17	10	15	19	13	7	4
peachpuff	10	27	29	19	19	7	22	18	10	8
peru	141	341	350	222	135	144	355	358	206	94
pink	56	96	85	64	36	30	54	54	38	16
plum	28	59	64	44	27	22	42	31	20	22
powderblue	245	532	390	141	75	136	270	205	87	45
purple	6	16	30	34	19	7	14	30	27	15
rebeccapurple	65	118	156	121	65	40	70	88	63	32
red	41	69	73	50	22	42	76	64	30	10
rosybrown	141	299	366	281	192	210	442	422	252	144
royalblue	47	83	62	32	13	33	48	40	16	10
saddlebrown	85	254	377	289	134	127	312	378	264	111
salmon	16	33	34	20	12	4	16	17	9	6
sandybrown	362	726	582	308	119	266	514	376	166	62
seagreen	119	215	204	97	34	98	182	170	86	26
seashell	70	148	129	67	38	39	91	113	132	148
sienna	148	356	418	306	153	119	310	389	264	116
silver	323	707	730	479	289	352	834	869	548	298
skyblue	499	932	668	310	152	282	525	366	188	83

slateblue	31	78	56	13	6	19	43	40	17	2
slategray	110	236	258	177	80	155	304	279	167	66
snow	824	1,725	1,377	751	501	690	1,397	1,096	572	363
springgreen	0	1	1	1	0	0	1	1	0	0
steelblue	421	845	710	352	138	358	660	537	268	89
tan	141	314	375	255	150	129	305	345	216	102
teal	196	337	309	209	118	160	314	287	169	75
thistle	47	95	108	82	47	41	78	79	63	28
tomato	125	236	203	122	65	77	134	105	72	43
turquoise	10	16	19	21	12	13	13	10	10	3
violet	1	5	9	8	5	1	3	6	6	2
wheat	164	424	445	232	136	119	322	333	195	106
white	2,459	4,024	2,737	1,419	1,188	972	1,639	1,235	807	661
whitesmoke	853	1,969	1,632	863	610	843	1,864	1,486	678	400
yellow	129	211	166	91	49	78	166	132	49	25
yellowgreen	186	360	335	198	81	146	247	202	116	44
Total	21,351	44,966	43,575	30,801	21,639	16,830	36,556	35,762	24,909	16,484

B Visual attributes

The following tables contain the results from the Turkey HSD tests and should be interpreted as follows: The "meandiff" column shows the difference in means between the two groups being compared, and the "p-adj" column shows the adjusted p-value for that comparison. A p-value of less than 0.05 indicates a statistically significant difference, while a p-value of greater than 0.05 indicates no significant difference. The "lower" and "upper" columns give the lower and upper bounds of the 95% confidence interval for the difference in means, and the "reject" column indicates whether the null hypothesis (that there is no difference in means between the two groups) can be rejected based on the p-value.

Table 3: Brightness statistical test data.

		Amazon				stat.	Open Library				stat.
Group 1	Group 2	meandiff	p-adj	lower	upper	sign.	meandiff	p-adj	lower	upper	sign.
0-3	4-6	-0.0144	0.0	-0.0181	-0.0107	True	-0.0095	0.0	-0.0144	-0.0045	True
0-3	7-9	-0.0576	0.0	-0.0613	-0.0538	True	-0.0442	0.0	-0.0492	-0.0393	True
0-3	10-12	-0.1194	0.0	-0.1234	-0.1154	True	-0.0913	0.0	-0.0966	-0.0860	True
0-3	13+	-0.1579	0.0	-0.1622	-0.1536	True	-0.1165	0.0	-0.1223	-0.1107	True
4-6	7-9	-0.0432	0.0	-0.0462	-0.0402	True	-0.0347	0.0	-0.0387	-0.0308	True
4-6	10-12	-0.1050	0.0	-0.1083	-0.1017	True	-0.0818	0.0	-0.0861	-0.0774	True
4-6	13+	-0.1435	0.0	-0.1472	-0.1398	True	-0.1070	0.0	-0.1120	-0.1020	True
7-9	10-12	-0.0618	0.0	-0.0652	-0.0585	True	-0.0471	0.0	-0.0514	-0.0427	True
7-9	13+	-0.1003	0.0	-0.104	-0.0966	True	-0.0723	0.0	-0.0773	-0.0673	True
10-12	13+	-0.0385	0.0	-0.0424	-0.0345	True	-0.0252	0.0	-0.0305	-0.0199	True

Table 4: Colorfulness statistical test data.

		Amazon				stat.	Open Library				stat.
Group 1	Group 2	meandiff	p-adj	lower	upper	sign.	meandiff	p-adj	lower	upper	sign.
0-3	4-6	-0.0120	0.0	-0.0150	-0.0090	True	-0.0143	0.0	-0.0177	-0.0109	True
0-3	7-9	-0.0260	0.0	-0.0291	-0.0230	True	-0.0331	0.0	-0.0365	-0.0296	True
0-3	10-12	-0.0531	0.0	-0.0563	-0.0499	True	-0.0684	0.0	-0.0721	-0.0647	True
0-3	13+	-0.0883	0.0	-0.0918	-0.0848	True	-0.1077	0.0	-0.1117	-0.1037	True
4-6	7-9	-0.0141	0.0	-0.0165	-0.0116	True	-0.0188	0.0	-0.0215	-0.0161	True
4-6	10-12	-0.0411	0.0	-0.0438	-0.0384	True	-0.0541	0.0	-0.0571	-0.0511	True
4-6	13+	-0.0764	0.0	-0.0794	-0.0734	True	-0.0934	0.0	-0.0968	-0.0899	True
7-9	10-12	-0.0271	0.0	-0.0297	-0.0244	True	-0.0353	0.0	-0.0384	-0.0323	True
7-9	13+	-0.0623	0.0	-0.0653	-0.0593	True	-0.0746	0.0	-0.0781	-0.0711	True
10-12	13+	-0.0352	0.0	-0.0385	-0.0320	True	-0.0393	0.0	-0.0430	-0.0356	True

Table 5: Contrast statistical test data.

		Amazon				stat.	Open Library				stat.
Group 1	Group 2	meandiff	p-adj	lower	upper	sign.	meandiff	p-adj	lower	upper	sign.
0-3	4-6	0.0010	0.2661	-0.0004	0.0023	False	0.0003	0.9761	-0.0012	0.0019	False
0-3	7-9	0.0052	0.0	0.0038	0.0065	True	0.0022	0.0014	0.0006	0.0038	True
0-3	10-12	0.0108	0.0	0.0094	0.0122	True	0.0028	0.0001	0.0011	0.0045	True
0-3	13+	0.0173	0.0	0.0158	0.0189	True	0.0001	1.0	-0.0018	0.0020	False
4-6	7-9	0.0042	0.0	0.0031	0.0053	True	0.0019	0.0005	0.0006	0.0031	True
4-6	10-12	0.0098	0.0	0.0086	0.0110	True	0.0025	0.0	0.0011	0.0039	True
4-6	13+	0.0164	0.0	0.0150	0.0177	True	-0.0003	0.9916	-0.0019	0.0013	False
7-9	10-12	0.0056	0.0	0.0044	0.0068	True	0.0006	0.7602	-0.0008	0.0020	False
7-9	13+	0.0122	0.0	0.0108	0.0135	True	-0.0021	0.0027	-0.0037	-0.0005	True
10-12	13+	0.0066	0.0	0.0051	0.0080	True	-0.0028	0.0001	-0.0045	-0.0010	True

Table 6: Entropy statistical test data.

		Amazon				stat.	Open Library				stat.
Group 1	Group 2	meandiff	p-adj	lower	upper	sign.	meandiff	p-adj	lower	upper	sign.
0-3	4-6	0.0235	0.0	0.0205	0.0264	True	0.0120	0.0	0.0091	0.0149	True
0-3	7-9	0.0404	0.0	0.0374	0.0433	True	0.0165	0.0	0.0135	0.01940	True
0-3	10-12	0.0357	0.0	0.0326	0.0389	True	-0.0043	0.0014	-0.0074	-0.0012	True
0-3	13+	0.0122	0.0	0.0088	0.0156	True	-0.0412	0.0	-0.0446	-0.0378	True
4-6	7-9	0.0169	0.0	0.0145	0.0193	True	0.0044	0.0	0.0021	0.0067	True
4-6	10-12	0.0123	0.0	0.0096	0.0149	True	-0.0163	0.0	-0.0189	-0.0138	True
4-6	13+	-0.0113	0.0	-0.0142	-0.0083	True	-0.0532	0.0	-0.0561	-0.0503	True
7-9	10-12	-0.0046	0.0	-0.0073	-0.0020	True	-0.0208	0.0	-0.0233	-0.0182	True
7-9	13+	-0.0282	0.0	-0.0311	-0.0252	True	-0.0576	0.0	-0.0606	-0.0547	True
10-12	13+	-0.0235	0.0	-0.0267	-0.0204	True	-0.0369	0.0	-0.0400	-0.0337	True

C Object Detection

Table 7: Frequency data, top 10 objects per age group.

	Amazon					Open Library				
Objects	0-3	4-6	7-9	10-12	13+	0-3	4-6	7-9	10-12	13+
Object 1	animal	animal	footwear	woman	woman	animal	animal	footwear	woman	poster
Object 2	chicken	footwear	animal	footwear	girl	chicken	footwear	animal	footwear	person
Object 3	toy	chicken	tree	girl	lipstick	toy	chicken	woman	girl	woman
Object 4	rabbit	toy	woman	man	man	rabbit	toy	mammal	man	girl
Object 5	footwear	tree	mammal	tree	footwear	footwear	mammal	tree	book	man
Object 6	tree	rabbit	chicken	animal	hair	mammal	tree	chicken	animal	lipstick
Object 7	mammal	mammal	toy	lipstick	fashion accessory	tree	rabbit	toy	tree	footwear
Object 8	carnivore	hat	hat	hat	jeans	carnivore	carnivore	hat	mammal	book
Object 9	teddy bear	carnivore	girl	mammal	dress	teddy bear	hat	girl	lipstick	fasion accessory
Object 10	hat	flower	boy	hair	eye	frog	dinosaur	man	hat	dress

Table 8: Correlation data Amazon, 10 objects with most extreme correlation values.

	Amazon									
Objects	0-3		4-6		7-9		10-12		13+	
Max 5	rabbit	0.15	chicken	0.15	person	0.17	woman	0.11	woman	0.30
	chicken	0.12	animal	0.14	clothing	0.14	girl	0.10	lipstick	0.26
	toy	0.12	rabbit	0.14	face	0.09	man	0.09	girl	0.17
	teddy bear	0.10	toy	0.13	poster	0.08	face	0.08	man	0.15
	animal	0.10	pig	0.10	footwear	0.07	jeans	0.06	eye	0.15
Min 5	lipstick	-0.09	jeans	-0.10	hand	-0.05	teddy bear	-0.08	rabbit	-0.11
	face	-0.10	man	-0.13	hair	-0.06	animal	-0.09	toy	-0.13
	girl	-0.10	girl	-0.15	eye	-0.07	toy	-0.09	chicken	-0.13
	man	-0.10	lipstick	-0.15	woman	-0.11	rabbit	-0.11	animal	-0.14
	woman	-0.15	woman	-0.21	lipstick	-0.13	chicken	-0.11	person	-0.15

Table 9: Correlation data Open Library, 10 objects with most extreme correlation values.

	Open Library									
Objects	0-3		4-6		7-9		10-12		13+	
Max 5	poster	0.17	poster	0.33	poster	0.34	poster	0.22	woman	0.28
	rabbit	0.15	person	0.21	person	0.29	face	0.17	lipstick	0.24
	chicken	0.14	clothing	0.21	clothing	0.28	woman	0.14	man	0.15
	toy	0.14	chicken	0.17	face	0.22	clothing	0.14	girl	0.14
	animal	0.11	animal	0.17	footwear	0.11	person	0.12	eye	0.12
Min 5	jeans	-0.04	jeans	-0.06	cosmetics	-0.02	pig	-0.05	mammal	-0.06
	lipstick	-0.06	girl	-0.07	woman	-0.02	toy	-0.06	rabbit	-0.08
	girl	-0.06	man	-0.07	hand	-0.02	teddy bear	-0.06	toy	-0.09
	man	-0.07	lipstick	-0.09	eye	-0.04	chicken	-0.07	chicken	-0.09
	woman	-0.09	woman	-0.12	lipstick	-0.07	rabbit	-0.08	animal	-0.10

Table 10: Co-occurence data Amazon, top 5 pairs with highest PMI.

	Amazon				
Top 10 Pairs	0-3	4-6	7-9	10-12	13+
Pair 1	bathroom accessory toilet paper	bidet toilet	jacuzzi tap	food processor salt and pepper shakers	sink waste container
Pair 2	plumbing fixture toilet paper	cupboard gas stove	toilet bidet	printer treadmill	bell pepper orange
Pair 3	bathroom accessory plumbing fixture	bidet tap	bathroom accessory toilet	light switch telephone	oyster wine rack
Pair 4	bidet toilet paper	tap toilet	printer treadmill	common fig zucchini	oyster sushi
Pair 5	bidet bathroom accessory	shrimp sushi	jacuzzi sink	sink tap	ant squash

Table 11: Co-occurence data Open Library, top 5 pairs with highest PMI.

	Open Library				
Top 10 Pairs	0-3	4-6	7-9	10-12	13+
Pair 1	tap telephone	ipod remote control	bathroom accessory plumbing fixture	food processor salt and pepper shakers	food processor salt and pepper shakers
Pair 2	countertop gas stove	cupboard gas stove	bathroom accessory sink	gas stove wok	jacuzzi sink
Pair 3	telephone toilet paper	countertop cupboard	plumbing fixture sink	bidet toilet	jaguar leopard
Pair 4	tap toilet paper	medical equipment toothbrush	potato radish	gas stove pizza	racket tennis racket
Pair 5	mixer paper towel	dumbbell flashlight	cupboard gas stove	countertop kitchen knife	squash zucchini

D Implied Story

Table 12: Named Entity Recognition data, expressed in mentions per thousand books + statistical significance

	Amazon					stat.	Open Library					stat.
Labels	0-3	4-6	7-9	10-12	13+	sign.	0-3	4-6	7-9	10-12	13+	sign.
ANIMAL	603.60	583.59	414.91	258.27	128.38	True	580.59	520.46	403.05	247.11	127.03	True
CARDINAL	68.69	69.84	74.82	65.51	44.45	True	68.83	69.69	74.26	67.01	47.48	True
DATE	26.23	26.01	20.52	15.15	14.65	True	26.6	25.87	21.16	14.88	13.13	True
EVENT	0.05	0.09	0.23	0.23	0.05	False	0.06	0.05	0.17	0.16	0.0	False
FAC	0.19	0.38	0.57	0.68	0.41	True	0.12	0.33	0.36	0.32	0.24	False
GPE	16.79	16.45	14.94	13.25	8.63	True	15.7	13.42	13.66	13.44	6.95	True
LANGUAGE	0.14	0.15	0.11	0.06	0.05	False	0.12	0.11	0.11	0.08	0.06	False
LAW	0.0	0.0	0.0	0.0	0.0	False	0.0	0.0	0.0	0.0	0.0	False
LOC	0.42	0.49	0.68	0.81	0.41	False	0.3	0.55	0.75	0.76	0.48	False
MONEY	0.28	0.31	0.3	0.26	0.05	False	0.3	0.33	0.2	0.2	0.06	False
NORP	3.3	5.26	6.99	5.96	3.95	True	2.96	5.29	7.39	6.74	3.96	True
ORDINAL	5.49	4.31	3.83	4.19	3.54	True	5.92	4.37	3.65	4.27	3.54	True
ORG	21.44	23.12	22.96	19.41	13.22	True	21.86	21.45	21.32	18.91	13.85	True
PERCENT	0.05	0.02	0.02	0.03	0.0	False	0.06	0.03	0.03	0.04	0.0	False
PERSON	60.09	63.57	74.34	86.08	78.47	True	59.36	59.86	69.43	80.73	74.34	True
PRODUCT	0.28	0.51	0.75	0.77	1.06	True	0.36	0.57	0.78	0.52	0.54	False
QUANTITY	0.65	0.4	0.34	0.35	0.18	True	1.07	0.57	0.42	0.68	0.6	False
TIME	6.51	5.39	4.51	5.48	6.34	True	6.75	5.21	3.37	3.47	5.34	True
WORK_OF_ART	0.19	0.13	0.16	0.19	0.05	False	0.18	0.16	0.2	0.2	0.0	False

Table 13: Topic modeling data, top 10 words per topic.

	Top 10 words									
Topic 1	children	two	people	group	image	three	moon	playing	middle	summer
Topic 2	words	written	house	water	city	secret	wall	story	last	high
Topic 3	title	long	bear	snow	baby	life	flower	paper	sweet	running
Topic 4	woman	man	holding	painting	drawing	front	hand	bird	sign	child
Topic 5	horse	tree	riding	hands	night	grass	angel	person	adventures	christmas
Topic 6	girls	flying	sky	two	world	character	kissing	air	princess	street
Topic 7	standing	black	white	front	close	face	background	red	next	blue
Topic 8	girl	dog	cat	pink	little	dress	holding	car	dragon	road
Topic 9	sitting	top	person	women	couple	laying	next	table	heart	beach
Topic 10	boy	young	girl	window	looking	star	bunch	school	field	fire

Table 14: Topic modeling data, expressed in topics per thousand books.

Topics	Amazon 0-3	Amazon 4-6	Amazon 7-9	Amazon 10-12	Amazon 13+	stat. sign.	Open Library 0-3	Open Library 4-6	Open Library 7-9	Open Library 10-12	Open Library 13+	stat. sign.
Topic 1	176	183	164	108	75	True	210	215	182	113	78	True
Topic 2	64	66	78	100	115	True	63	65	77	98	109	True
Topic 3	96	74	59	59	58	True	93	73	60	68	79	True
Topic 4	104	118	133	158	182	True	104	117	131	150	172	True
Topic 5	70	73	76	74	65	True	66	71	76	77	68	True
Topic 6	71	70	70	62	50	True	71	70	69	62	50	True
Topic 7	97	94	106	157	226	True	87	87	105	159	225	True
Topic 8	150	142	123	92	58	True	148	139	124	95	61	True
Topic 9	75	77	87	100	111	True	67	71	81	93	98	True
Topic 10	97	101	104	92	60	True	91	93	94	85	60	True