

Covering Covers

Characterization Of Visual Elements Regarding Sleeves

Beyhan, Yessin; Pera, Maria Soledad

DOI

[10.1145/3563359.3597404](https://doi.org/10.1145/3563359.3597404)

Publication date

2023

Document Version

Final published version

Published in

UMAP 2023 - Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization

Citation (APA)

Beyhan, Y., & Pera, M. S. (2023). Covering Covers: Characterization Of Visual Elements Regarding Sleeves. In *UMAP 2023 - Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization* (pp. 28-33). (UMAP 2023 - Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3563359.3597404>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Covering Covers: Characterization Of Visual Elements Regarding Sleeves

Yessin Beyhan
yessinbeyhan@gmail.com

Web Information Systems - TU Delft
The Netherlands

Maria Soledad Pera
m.s.pera@tudelft.nl

Web Information Systems - TU Delft
The Netherlands

ABSTRACT

The aim of this work is to explore common traits preferred across different age groups of children to identify the appeal of book covers. By analyzing visual attributes, visible objects, and implied stories inferred from the covers, we can gain insights into the elements that are most attractive to children up to 18 years old. These findings can then contribute towards advancing personalization for recommender systems for children through new means that do not rely on historical data, seldom available for this user group.

CCS CONCEPTS

• **Social and professional topics** → **Children**; • **Information systems** → *Personalization*.

KEYWORDS

Recommender Systems, user modeling, personalization, children

ACM Reference Format:

Yessin Beyhan and Maria Soledad Pera. 2023. Covering Covers: Characterization Of Visual Elements Regarding Sleeves. In *UMAP '23 Adjunct: Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (UMAP '23 Adjunct)*, June 26–29, 2023, Limassol, Cyprus. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3563359.3597404>

1 INTRODUCTION

Encouraging children's reading habits is imperative for their cognitive development [3] and for shaping their interests and values [27]. However, selecting the right book to foster these habits for a child can be challenging due to many reasons, including the vast number of options available [24], and the wide range of reading levels [12] and interests among young readers [2, 23]. Recommender systems (RS) could ease this task by offering appealing and suitable reading materials [26]. Still, research focused on improving book RS as well as popular RS available via common online platforms like Amazon have mostly served adult populations [4]. In fact, to the best of our knowledge, literature on RS that specifically cater to children, particularly in the book domain, is very limited [5, 25, 26]. This is primarily attributed to limited children's user-system interactions that could drive the design and enable assessment of such RS [7, 8].

Focused on the book domain, we argue that in order to improve the design of RS for children, it is necessary to understand the

factors that contribute to their appeal. Book covers play a crucial role in attracting the attention of young readers [15, 29]. Previous research has shown that book covers can impact the book's appeal to children of different ages [13]. While existing RS like Storytime [20] have explored the use of images to elicit preferences, there is still limited understanding of the specific factors that contribute to the appeal of children's book covers and how these factors may differ as children mature. More so, existing book RS for children [5, 25, 26] have yet to leverage cover information into their recommendation algorithms. As a step towards advancing knowledge in this area, we analyze the impact on children's book cover appeal of three specific aspects of covers: *visual attributes*, *visible objects*, and *implied story*, in relation to different age groups. Specifically, 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, 10, 33]. For our initial exploration, we rely on 93,662 books from the Goodreads [36, 37] and Open Library [17] datasets, which provide metadata and cover images on 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.

With this work, we provide insights that can ultimately influence the design of children's book RS by identifying key factors that contribute to book cover appeal for different age groups. Lessons learned yield representative heuristics that can facilitate the representation of young users, leading to performance improvements in terms of accuracy and, more importantly, personalization while respecting their privacy and avoiding intrusive data collection [30]. We posit that this preliminary study can help lay the groundwork for more effective and responsible RS for children's books, ultimately promoting a lifelong love of reading. In the interest of transparency and reproducibility, we make the scripts used for this work publicly available: <https://github.com/Yessin111/Childrens-Book-Cover-Analysis>.

2 EXPERIMENTAL SET UP

In this section, we describe the user groups, data, and perspectives considered in the experiments conducted for exploration purposes.

Age groups. Inspired by the empirical explorations presented in [19, 32], we consider the appeal of covers among children of different ages, as similarities and differences are likely to be more pronounced than those between individual ages. We define 5 age groups 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 and language development [10].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UMAP '23 Adjunct, June 26–29, 2023, Limassol, Cyprus

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9891-6/23/06.

<https://doi.org/10.1145/3563359.3597404>

Children at *ages 4-6* focus on developing basic language skills [33]. 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 [33]. At *ages 10-12*, the focus shifts to utilizing the information and knowledge gained from reading texts and books in other contexts [33]. Young adult literature is typically intended for readers aged *13-15 and older* and is characterized by themes relevant to adolescents, such as relationships, identity, and finding one's place in the world [1].

Data description. The data considered for this exploration consists of 93,662 unique children's books from various English publishing authors, spanning a range of reading levels from early childhood to young adult (see Table 1). For this, we turn to the Goodreads dataset (GR) [36, 37], specifically books¹ in English from the "children" and "young adult" datasets, with each including metadata like the title, author, country of publishing, age, and cover. Aware that book covers can vary across editions of the same book, whenever possible, we also examine book covers from Open Library [17].

Table 1: Distribution of books across age groups. Note that in most cases, books are associated with an age-range, not a singular age/grade and therefore can be associated with more than one of the age groups we consider in our exploration.

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

Perspectives of exploration. Below we briefly present the aspects considered for analysis purposes.

Visual attributes. We consider five lenses for this aspect: dominant color, brightness, colorfulness, contrast, and entropy [19]. To determine the dominant color of each cover image, we use k-means clustering with $k=9$, to account for the 7 colors of the rainbow plus black and white. To measure the brightness, we convert the images to grayscale and take the average pixel brightness. The colorfulness is calculated using the colorfulness metric methodology described in [9]. For finding the contrast, we take the Root Mean Square contrast definition. The entropy is calculated using Shannon's entropy [31]. Entropy, in the context of cover images, is a measure of the disorder of the colors in the image. We normalize all of these calculated visual attributes to a range of [0,1] to make it easier to analyze the values.

Object detection. Objects within 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. However, identifying objects is not a trivial task, as detection algorithms are typically trained on a small number of detection classes such as animals, vehicles, or people [18]. To identify objects on covers, we employ the OWL-ViT detection model [21], a zero-shot classification model, together with 545 labels from Google's ImageNet dataset [14], which

includes a wide range of object classes. This allows for the identification of common objects with a probability in a range of [0,1], without the need for training data on those objects².

Implied story. The analysis of the story conveyed by a book cover goes beyond just identifying objects, as it also examines characters and emotions. For this, we use the BLIP captioning model [16], a state-of-the-art machine learning tool that generates descriptive captions for images. One of the strengths of such a 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.

To *illustrate* the type of information we infer from book covers to inform our exploration, consider the cover of the book "Lost and Found" by Oliver Jeffers (extracted from the GR dataset), intended for children aged 0-3 in Figure 1.

The values for the extracted *visual attributes* are: (i) Dominant Color: (13, 172, 226), (ii) Dominant Color Name: blue, (iii) Brightness: 0.615, (iv) Colorfulness: 0.519, (v) Contrast: 0.171, (vi) Entropy: 0.955.

The dominant color –blue– is fitting as the sea is the largest 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.

The *objects* extracted using the object detection are: (i) Poster: 0.617, (ii) Penguin: 0.485, (iii) Clothing: 0.442, (iv) Person: 0.428, (v) Hat: 0.274, (vi) Paddle: 0.267, (vii) Canoe: 0.24, (viii) 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.

The emerging *implied story* is: "A picture of a penguin and a man in a boat". The caption –which captures the essence of the image




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

¹Books that fall into more than two age groups are deemed as having too broad of a range and are excluded from this exploration.

²Using sample covers, we set a threshold for certainty=20%, i.e., we treat as objects those that the model is able to detect with a probability of 0.2 or higher.

and its contents– identifies the key elements present in the image. It is worth noting that, again, 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.

This example evinces the potential of looking into covers from different perspectives in order to discern aspects that can inform the appeal of book covers for children in varied age groups.

Analysis. We conduct several investigations in order to explore the prominence, or lack thereof, of certain visual elements present in covers of books intended for specific user groups. Specifically, we juxtapose the results obtained for the aspects associated with each of the aforementioned perspectives of exploration.

For the trends of dominant color usage, we compute and compare the proportion of all dominant colors present per age group. For the remaining visual attributes, we examine any significant differences in their values across the different age groups and sources using a one-way ANOVA with $p < 0.05$ to determine whether there is a significant difference in the means of the attribute values for each age group and then using a Turkey HSD test to identify which specific pairs of age groups show significant differences.

In the case of object detection, we seek to understand the prevalence and relationship between objects found on children’s book covers across different age groups: *frequency analysis*, which shows us which objects are more prominent in each age group, *correlation analysis* using Pearson correlation, which enables us to identify which objects are indicative of each age group, and *co-occurrence analysis* using point-wise mutual information, which identifies which objects frequently appear together on covers within the same age group. Lastly, to analyze the content of the captions in the children’s books, we use *Named Entity Recognition* (NER) and *Topic Modeling*. The former is a method for extracting information and identifying entities such as people, organizations, and locations from unstructured text. The latter is a technique for identifying underlying topics or themes within a corpus of text. We used the spaCy [11] library for NER and Gensim [28] library for Topic Modeling. We also report on statistically significant shifts in patterns or themes present in the text data across different age groups using chi-squared tests with $p < 0.05$, which are able to quantitatively measure this association.

3 RESULTS AND DISCUSSION

Here, we present the outcomes emerging from the conducted empirical investigations. Unless mentioned otherwise, significant trends remain, regardless of the source of the covers considered.

3.1 Visual attributes

We discuss salient findings pertaining to visual attributes. We illustrate results for dominant colors in Figure 2; we capture results for the remaining attributes in Figure 3.

Dominant color. We first consider the distribution of dominant colors for the cover images. White and black were the most common colors, making up 50-70% of all covers. As age increases, the percentage of white decreases and the percentage of other colors tend to increase. The shift was statistically significant for every

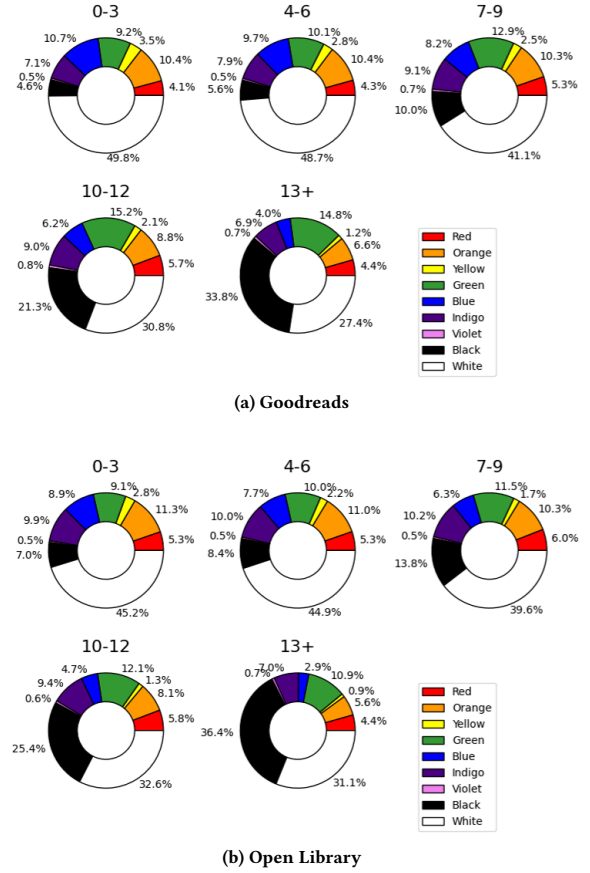


Figure 2: Dominant colors of cover images per age group.

dominant color in each age group. Orange, yellow, and blue percentages decrease as age increases, while green and violet percentages increase. 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 red and indigo may be particularly appealing to children in the middle age groups.

Brightness. There are significant differences in brightness between all age groups, with each group being overall darker than its predecessor. We posit that this could be due to the assumption that older children relate more to mature and serious themes [34], which are often conveyed through darker colors.

Colorfulness. The colorfulness of book covers decreases with statistical significance 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. This could reflect a preference for more mature and realistic themes.

Contrast. The contrast attribute shows significant differences in book covers across all age groups, except for ages 0-3 and ages 4-6 for both Goodreads and Open Library, and ages 0-3 and 13+, 4-6 and 13+, and 7-9 and 10-12 for Open Library. There may be a trend towards increasing contrast in book covers as the age group increases. Further research is needed to confirm this trend and to

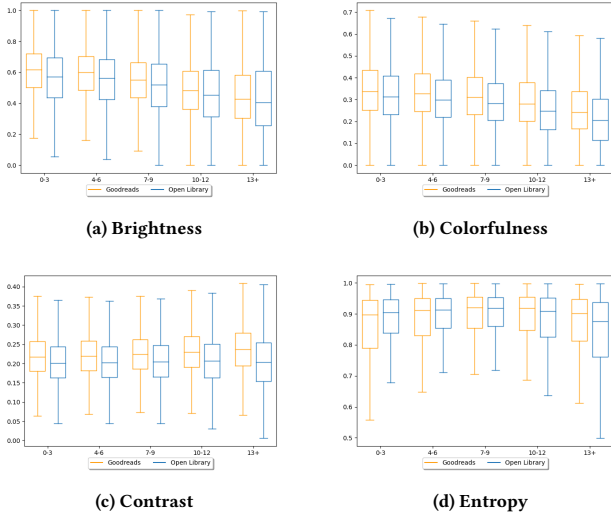


Figure 3: Visual attributes of cover images per age group.

determine the precise nature of the relationship between contrast and age group.

Entropy. The entropy values among different age groups were statistically significant, with an increase in entropy from ages 0-3 to ages 7-9 and a 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.

3.2 Object detection

We discuss findings arising from frequency and correlation analysis of the objects detected. Due to the large number of labels used for the zero-shot detection, we focus on discussing the most notable results.

Frequency. The most prominent objects detected on the book covers across all age groups and sources were “poster”, “clothing”, “person”, and “face”. The prevalence of objects associated with people indicates that human figures are common elements on children’s book covers. In the younger age groups (0-3 and 4-6), nature and toys were commonly represented. However, as the age group increases, there is a shift towards people-related objects and personal accessories. The trend continues in the 10-12 and 13+ age groups, where there is a stronger focus on people and their personal features and belongings. This suggests that illustrations of animals and nature are more common in books aimed at younger readers, while books aimed at older readers are more likely to feature illustrations of people.

Correlation. Although correlation analysis showed weak correlations between the frequency of objects on book covers and age groups, the objects observed in the analysis align with the patterns seen in the frequency analysis. For the youngest age group, objects related to animals and toys had a positive correlation, while for the oldest age groups, objects related to people had a positive correlation. The negative correlation for both ends of the age groups was

inverse to each other, indicating 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. Therefore, while objects on covers cannot be used to accurately predict the targeted age group, the objects present align with the observed patterns in age group preferences.

Co-occurrence. The study found some trends in the types of objects commonly found together on book covers for different age groups, but these do not provide clear insights into the appeal of books to different age groups. For instance, we see a shift from toilet-related objects to food and kitchen-related objects as age increases, but these results are anecdotal, i.e., not in line with the findings from the frequency and correlation analysis.

3.3 Implied story

We discuss salient findings resulting from NER and topic modeling analysis on yielded implied stories.

Named Entity Recognition. Outcomes from NER analysis show significant differences in the frequency of mention of various entities across age groups. Labels such as animals, numbers, dates, locations, nationalities, organizations, people, and time were mentioned less as age increased, while events, languages, laws, locations, money, percentages, and works of art did not vary significantly with age. These trends suggest that books aimed at older age groups may be more focused on serious topics and have less visually rich covers. These findings were consistent across both datasets, indicating that they may be generalizable to other book cover captions as well.

Topic Modeling. We examined 10 distinct topics computed on all books over all age groups, each associated with a set of words that suggest the subject matter of the topic. We found that there was a significant variation in the distribution of captions across the different age groups. For example, the topic associated with the words “children”, “playing”, and “summer” 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, the topic associated with the words “dog”, “cat” and “dragon”, but also “girl”, “pink”, “little”, and “dress” 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.

The topic associated with the words “beach”, “couple” and “heart” 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. The topic associated with the words “young”, “boy”, “school” and “field” is more prevalent in the 7-9, 10-12 and 13+ age groups, with a decrease in frequency as age decreases. These findings are anticipated, 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, offering insights into how the themes and content of the covers change with the intended audience.

4 CONCLUDING REMARKS, LIMITATIONS, AND FUTURE WORK

With this work, we aimed to understand the appeal of book covers to children of different ages. We focused on three aspects of covers: their visual attributes, depicted objects, and 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 is 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. We will consider collecting and labeling a large dataset of children's book covers to train a more specialized object detection algorithm to mitigate any errors caused by the model.

Traditionally, RS leverage content-based features, such as readability levels, topics, or ratings, or historical user-system interactions to provide recommendations. We argue that the results of this initial exploration can have direct implications for the design of RS, in terms of leveraging reported trends and patterns as part of new recommendation algorithms that could then produce more personalized book recommendations for children without requiring the often historical data in the form of ratings and interactions required to train and deploy RS.

By taking into account the age-appropriate visual attributes of book covers, such as brightness and colorfulness, a RS can provide recommendations that match the preferences of younger readers. Similarly, understanding which objects on book covers are more appealing to certain age groups can help prioritize books with those objects for corresponding age groups, resulting in more accurate and effective book recommendations. Moreover, incorporating visual features into RS could contribute towards addressing the cold start problem, where the system struggles to provide accurate recommendations for new users or items. By recommending books based on visual attributes, a RS can provide personalized suggestions even for new users or items. Lessons learned discussed in this work also highlight the need to balance personalization and privacy. By leveraging visual features of book covers, a RS can provide personalized book recommendations without collecting sensitive user data. This can lead to a more privacy-preserving RS that still provides an accurate and personalized experience for users.

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 [22] 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 [38] or even be applied to completely different areas of recommendation, such as analyzing the appeal of movie [6] or music album covers [35]. 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.

REFERENCES

- [1] Linda Ackermans. 2021. Een decennium young adult literatuur in Nederland (2009-2019): een stand van zaken. *Tijdschrift voor Nederlandse Taal- en Letterkunde* 137, 1 (2021), 45–72.
- [2] Richard E Ahl and Yarrow Dunham. 2020. Have your cake, and your asparagus, too: Young children expect variety-seeking behavior from agents with diverse desires. *Cognitive Development* 54 (2020), 100882.
- [3] Richard C Anderson, Larry L Shirey, Paul T Wilson, and Linda G Fielding. 2021. Interestingness of children's reading material. In *Aptitude, learning, and instruction*. Routledge, 287–300.
- [4] Khalid Anwar, Jamshed Siddiqui, and Shahab Saquib Sohail. 2019. Machine learning techniques for book recommendation: an overview. In *Proceedings of International Conference on Sustainable Computing in Science, Technology and Management (SUSCOM)*, Amity University Rajasthan, Jaipur-India.
- [5] Yufeng Chen, Jun Ma, Yaqiong Du, Rui Sun, and Bojuan Niu. 2021. Research on Elementary School Students' Books Recommendation Algorithm Based on Words and Character Library. In *2021 2nd International Conference on Artificial Intelligence and Information Systems*. 1–5.
- [6] Cecilia Fagerholm. 2009. The use of colour in movie poster design: An analysis of four genres. *N.A.* (2009).
- [7] Jerry Alan Fails, Maria Soledad Pera, Franca Garzotto, and Mirko Gelsomini. 2017. 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*. 376–377.
- [8] Emilia Gómez, Vicky Charisi, and Stephane Chaudron. 2021. Evaluating Recommender Systems with and for Children: towards a Multi-Perspective Framework.. In *Perspectives@ RecSys*.
- [9] David Hasler and Sabine E Suesstrunk. 2003. Measuring colorfulness in natural images. In *Human vision and electronic imaging VIII*, Vol. 5007. SPIE, 87–95.
- [10] Elizabeth A Hasson. 1991. "Reading" with infants and toddlers. *Day Care and Early Education* 19, 1 (1991), 35–37.
- [11] Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. (2017). To appear.
- [12] Charles Hulme and Margaret J Snowling. 2008. *Developmental disorders of language learning and cognition*. Wiley-Blackwell, Chichester, England.
- [13] Leigh Ann Jones. 2006. *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.
- [14] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. 2020. The open images dataset v4. *International Journal of Computer Vision* 128, 7 (2020), 1956–1981.
- [15] Monica Landoni and Elisa Rubegni. 2017. Aesthetic relevance when selecting multimedia stories. *International Workshop on Children & Recommender Systems* (2017).
- [16] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086* (2022).
- [17] Open Library. 2015. Open Library. <https://openlibrary.org/>.
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [19] Ashlee Milton, Leveson Batista, Garrett Allen, Siqi Gao, Yiu-Kai D Ng, and Maria Soledad Pera. 2020. "Don't judge a book by its cover": Exploring book traits children favor. In *Fourteenth ACM Conference on Recommender Systems*. 669–674.
- [20] Ashlee Milton, Michael Green, Adam Keener, Joshua Ames, Michael D Ekstrand, and Maria Soledad Pera. 2019. StoryTime: eliciting preferences from children for book recommendations. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 544–545.
- [21] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. 2022. Simple Open-Vocabulary Object Detection with Vision Transformers. *arXiv preprint arXiv:2205.06230* (2022).
- [22] Taiga Miyazono, Brian Kenji Iwana, Daichi Haraguchi, and Seichi Uchida. 2021. Font Style that Fits an Image—Font Generation Based on Image Context. In

- International Conference on Document Analysis and Recognition*. Springer, 569–584.
- [23] Kathleen AJ Mohr. 2006. Children's choices for recreational reading: A three-part investigation of selection preferences, rationales, and processes. *Journal of Literacy Research* 38, 1 (2006), 81–104.
 - [24] N.A. 2022. *LIBRARY AND BOOK TRADE ALMANAC 2022: Formerly the Bowker annual*. N.A.
 - [25] Yiu-Kai Ng. 2020. CBRec: a book recommendation system for children using the matrix factorisation and content-based filtering approaches. *International Journal of Business Intelligence and Data Mining* 16, 2 (2020), 129–149.
 - [26] Maria Soledad Pera and Yiu-Kai Ng. 2014. Automating readers' advisory to make book recommendations for k-12 readers. In *Proceedings of the 8th ACM Conference on Recommender Systems*. 9–16.
 - [27] Husni Rahim and Maila Dinia Husni Rahiem. 2012. The use of stories as moral education for young children. *International Journal of Social Science and Humanity* 2, 6 (2012), 454.
 - [28] Radim Rehurek and Petr Sojka. 2011. Gensim–python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic* 3, 2 (2011).
 - [29] Kara Reuter. 2007. Assessing aesthetic relevance: Children's book selection in a digital library. *Journal of the American Society for Information Science and Technology* 58, 12 (2007), 1745–1763.
 - [30] Elaine Rich. 1979. User modeling via stereotypes. *Cognitive science* 3, 4 (1979), 329–354.
 - [31] Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal* 27, 3 (1948), 379–423.
 - [32] Lawrence Spear, Ashlee Milton, Garrett Allen, Amifa Raj, Michael Green, Michael D Ekstrand, and Maria Soledad Pera. 2021. Baby Shark to Barracuda: Analyzing Children's Music Listening Behavior. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 639–644.
 - [33] Paul Van den Broek, Jolien M Mouw, and Astrid Kraal. 2015. Individual differences in reading comprehension. In *Handbook of individual differences in reading*. Routledge, 156–168.
 - [34] Helma van Lierop-Debrauwer and Neel Bastiaansen-Harks. 2005. *Over grenzen: De adolescentenroman in het literatuuronderwijs*. Vol. 6. Eburon Uitgeverij BV.
 - [35] Tara Venkatesan, Qian Janice Wang, and Charles Spence. 2020. Does the typeface on album cover influence expectations and perception of music? *Psychology of Aesthetics, Creativity, and the Arts* (2020).
 - [36] Mengting Wan and Julian McAuley. 2018. Item recommendation on monotonic behavior chains. In *Proceedings of the 12th ACM conference on recommender systems*. 86–94.
 - [37] Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian McAuley. 2019. Fine-grained spoiler detection from large-scale review corpora. *arXiv preprint arXiv:1905.13416* (2019).
 - [38] Alexis Weedon. 2016. In real life: book covers in the Internet bookstore. In *Judging a Book by Its Cover*. Routledge, 139–148.