RESEARCH BRIEF

What We Teach About Race and Gender: Representation in Images and Text of Children's Books

Based on BFI Working Paper 2021-44, "What We Teach About Race and Gender: Representation in Images and Text of Children's Books," by Anjali Adukia, Assistant Professor, UChicago's Harris School of Public Policy; Alex Eble, Assistant Professor, Teachers College Columbia University; Emileigh Harrison, PhD Student, UChicago's Harris School of Public Policy; Hakizumwami Birali Runesha, AVP and Director, UChicago's Research Computing Center; and Teodora Szasz, UChicago Computational Scientist.

KEY TAKEAWAYS

- ✓ Books are important for teaching children about the roles that they and others can or cannot inhabit.
- ✓ Given persistent racial and gender inequality in society, representations in books can offer a key means to address, perpetuate, or entrench core societal inequalities.
- ✓ To address these important questions about race and gender, this new research applies innovative techniques in AI to analyze images and text in children's books over time.
- ✓ A key finding of this research reveals that despite growing awareness in recent decades about race and gender issues in curricula, children's books generally skew toward lighter skin and male representation.

Educators and caregivers are generally thoughtful about choosing books to read to their young children, or when selecting books for children to read themselves. They may look for books that entertain, educate, and otherwise incorporate values that they hold dear. However, if those values include race and gender diversity, they will have to search a little harder.

New research employing path-breaking artificial intelligence (AI) tools reveals that characters in children's books, as measured by illustrations and text, are largely white and male. In "What We Teach About Race and Gender: Representation in Images and Text of Children's Books,' the authors find that this white/male dominance is even true of books published in recent decades during a period of heightened awareness about race and gender issues.

This research has important implications for educators and publishers, and others concerned about the influence of books on childhood development. In addition, the authors' novel methodology offers the promise of innovative investigations into other forms of text and visual media, including all types of literature and nonfiction, journalism, websites, art, photography, television, videos, movies, and many others.

Becker Friedman Institute
FOR ECONOMICS AT UCHICAGO

5757 S. University Ave.
Chicago, IL 60637
Main: 773.702.5599
bfi.uchicago.edu

A Note on Methodology

The authors' main data set is a series of books targeted to children and likely to appear in homes, classrooms, and school libraries over the past century. Specifically, they use books that received awards either administered or featured by the Association for Library Service to Children, a division of the American Library Association, starting in 1922. These and other children's books are often filled with images that transmit implicit and explicit messages to readers.

Historically, human coders provided content analysis, a time-consuming effort necessarily limited in scope and impacted by human behavior and biases. To address these limitations, the authors devised a unique application to harness the analytic power of artificial intelligence tools (AI). They developed computer vision tools that use convolutional neural networks to identify and classify components of images; in this case, detecting characters in photos and illustrations and classifying their race, gender, and age. While AI tools also reflect bias in their training data and algorithms, they can be more replicable, can be standardized, and can be applied to a much larger sample than manual content analysis.

Analyzing images involves three primary components: training the computer to detect faces, classifying skin color, and predicting the race, gender, and age of the faces. The authors build on existing face analysis software tools and also make pathbreaking improvements, including training their model to analyze illustrations, developing a classification of skin color, and introducing higher precision for classification of gender and age. [See the full working paper for detailed description and for many visual representations of this work, as well as an interactive chart depicting race and gender representation in various book collections.]

Diversity is not yet mainstream

Research has revealed the importance of curricular materials in education for teaching children about the world. In particular, the way that people are represented within books—including the roles that they inhabit—can contribute to children's understanding about what roles they and others can or cannot inhabit. Given persistent racial and gender inequality in society and the importance of identity and representation in driving beliefs, aspirations, academic effort, and outcomes, these representations offer a key means to either address, perpetuate, or entrench core societal inequalities.

So how do children's books stack up in terms of issues pertaining to race and gender? To answer this and related questions, the authors developed new software for the systematic analysis of images, highlighting their potential use in a wide range of applications in policy, education practice, and social science research. [See A Note on Methodology.] They then applied those tools, alongside established text analysis methods, to analyze children's books categorized broadly as Mainstream, or those considered of high literary value but written without explicit intention to highlight an identity group (e.g., the Newbery and Caldecott Awards); and Diversity books selected because they highlight experiences of specific identity groups (e.g., the Coretta Scott King and South Asia Book Awards).

In total, the 1,130 books in the study came from 19 different award categories and included over 160,000 pages of content published over the last 100 years (The Newbery was first awarded in 1922 and the Caldecott in 1938, for example, while the

> Books selected to highlight people of color or females increasingly depict characters with darker skin tones over time. However, Mainstream books are more likely to depict lighter-skinned characters than those in the Diversity collection, even conditional on character race.

Figure 1 · Image Results: Skin Color and Race Predictions

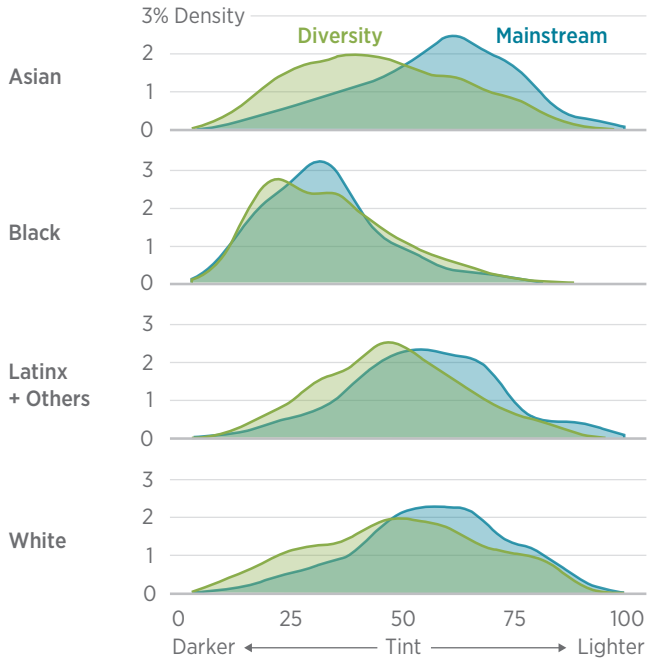

Note: This figure shows the distribution of skin color tint by predicted race of the detected faces in images. For example, the top graph in Figure 1 shows that a person identified as Asian in the Mainstream collection will be depicted with lighter skin on average than a person identified as Asian in the Diversity collection.

Coretta Scott King Awards began in 1970, and the South Asia Book Awards in 2012).

The authors' novel analysis of images revealed the following about race in children's books:

- Books in the Mainstream collection are more likely to depict lighter-skinned characters than those in the Diversity collection, even conditional on identified character race, potentially appealing to the assumed preferences of the median reader.
- Books selected to highlight people of color or females increasingly depict characters with darker skin tones over time. However, Mainstream books have increased representation of lighter skin tones over the last two decades despite growing rhetoric about the importance of diverse representation. Black and Latinx people are underrepresented in the images and text, relative to their share of the US population.
- Also, while females have always appeared in pictures over time (still less than 50 percent on average, but closer to 50 percent than in text), they are predominantly White females.
- Particularly surprising is that despite no systematic differences in skin tones across ages in society, children are more likely than adults to be shown with lighter skin, regardless of collection.

The authors also compared the incidence of female appearances in images to female mentions in text to find that:

- Females are more consistently visualized (seen) in images than spoken about (heard) in the text, except in the collection of books specifically selected to highlight females, suggesting symbolic inclusion of females in pictures without their substantive inclusion in the actual story.
- This underrepresentation holds regardless of the measure used: predicted gender of the pictured character, pronoun counts, specific gendered words, famous figure gender, and character first names.
- Males, especially White males, are persistently more likely to be represented by every measure, with little change over time despite substantial changes in female societal participation.
- Even though these books are targeted to children, adults are depicted more often than children in both images and text.

In a separate analysis about the appearance of famous figures, the authors find that:

- The Diversity collection has broader geographic representation of famous figures born outside of the United States or Europe than the Mainstream collection. However, when either collection presents a character outside of these two regions, that character is more likely to be male.
- This finding suggests that while the Diversity collection may represent a broader range of nationalities, it is still unequal in its representation of identity at the intersection of gender and nationality.
- Moreover, White males comprise the majority of famous figures in all collections. Famous people from other racial groups are less likely

Figure 2 · Women Should be Seen More Than Heard?

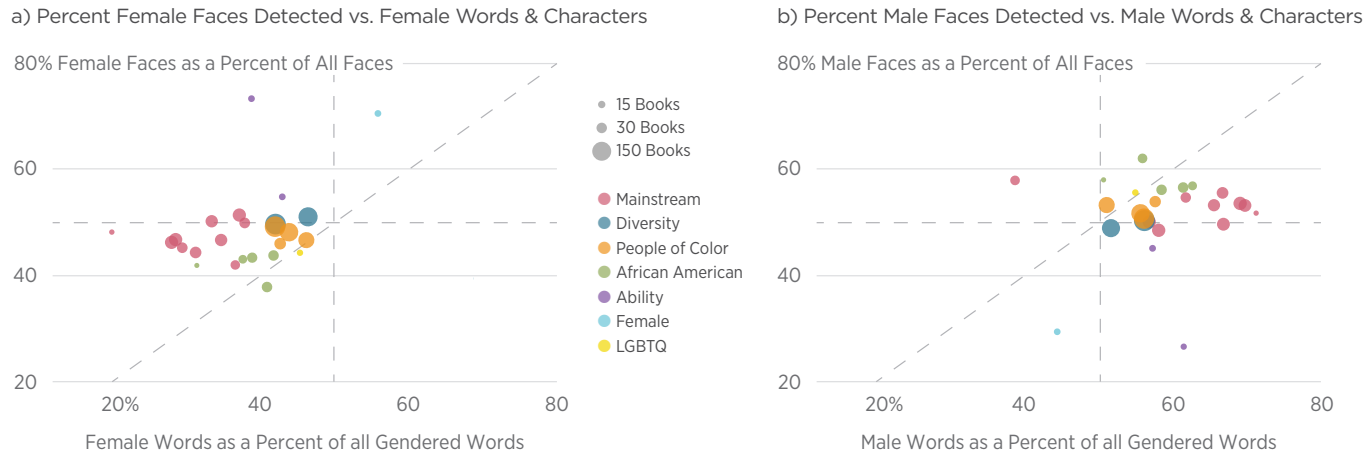

Note: In this figure the authors contrast the representation of females in the text of these collections of books with representation of females in the images of the same books. In Panel A, they plot collection-by-decade averages of female representation in images (on the y axis) and female representation in text (on the x axis). On the y-axis, the authors plot the average percent of female faces out of all faces detected. On the x-axis, they plot the average percent of gendered words which are female. Panel B shows the inverse to Panel A: the proportional representation of males in images and text. The authors detect faces using a Google Vision AutoML model trained on illustrations. Within these faces, the authors classify gender using an AutoML algorithm trained using a manually labeled random sample of our data, assigning the female value to all faces receiving the female label with a prediction value of greater than 50 percent.

than either White people or Black people to be represented in any collection (0 – 8 percent), but even then, males are generally more likely to be represented than females within any racial group.

The authors focus their analytical lens on a number of other questions, including the representation of characters from their place of origin (Mainstream books mostly feature people from Europe and the eastern United States, while Diversity books feature those from across the world, especially including the Southern Hemisphere), and the intersectionality of two or more characteristics in one figure, (for example, there is relatively low representation of Black women, even in books in the Diversity collections). [See the full working paper for detailed description and for many visual representations of this work, as well as an interactive chart depicting race and gender representation in various book collections.]

There are limitations to this analysis, as the authors stress, including the inherent subjectivity that

> Mainstream books consistently depict people within each race as having lighter skin than people of the same race books selected to highlight people of color.

can be built into AI tools, the flaws in current measures of gender identity, and algorithms' inability to perfectly detect faces or isolate skin from faces and thus leading to measurement error. Also, this analysis consists of a numerical accounting of different characters through simple representational statistics, that is, whether characters are included. However, if a character is depicted in a reductive or stereotypical manner, then solely the existence of representation will be insufficient and possibly counterproductive.

Conclusion

While many educators and schools wish to eliminate books that have overt racial and gender bias, such efforts are necessarily piecemeal and the judgments behind them subjective. This novel research program takes the adage "a picture is worth a thousand words" to heart and introduces a high level of objectivity by systematically analyzing images and text in prominent children's books. In very broad sum, despite growing awareness in recent decades, children's books generally skew toward lighter skin and male representation.

The authors are careful to stress that their work is not an attempt to offer a prescription for children's books or a scorecard for publishers. What is the "optimal" level of representation in children's books? That is a great question which is beyond the scope of this research and best left for experts in the field

CLOSING TAKEAWAY

Females are more consistently visualized (seen) in images than spoken about (heard) in the text, except in the collection of books specifically selected to highlight females, suggesting symbolic inclusion of females in pictures without their substantive inclusion in the actual story.

of education. That said, by offering a method to measure representation, this research does offer opportunities to address important issues and to better achieve desired goals.

Finally, the authors' innovative application of AI should lead to further development of tools that can measure how people are represented in books and other media, and thereby help determine what content depicts characters in their full humanity. A systemic problem requires a systemic solution. This work can stimulate a wide range of social science research that uses printed content—both images and text—as primary source data that can help us understand how variation in representation shapes human beliefs, behavior, and outcomes. It's a tall order. However, providing research that expands our understanding about diversity in content can help us overcome the structural inequality that pervades society and our daily lives.

READ THE WORKING PAPER

NO. 2021-44 · APRIL 2021

What We Teach About Race and Gender: Representation in Images and Text of Children's Books

bfi.uchicago.edu/working-paper/2021-44

ABOUT OUR SCHOLARS

Anjali Adukia
Assistant Professor, Harris School of Public Policy
harris.uchicago.edu/directory/anjali-adukia

Emileigh Harrison
PhD Student, Harris School of Public Policy
harris.uchicago.edu/news-events/news/student-profile-emileigh-harrison-phd-student

THE UNIVERSITY OF CHICAGO Harris Public Policy

Hakizumwami Birali Runesha
AVP and Director, UChicago's Research Computing Center
researchinnovation.uchicago.edu/leadership/h-birali-runesha/

Teodora Szasz
UChicago Computational Scientist
rcc.uchicago.edu/about-rcc/our-team

THE UNIVERSITY OF CHICAGO
Research, Innovation and National Laboratories
Research Computing Center

Written by David Fettig, BFI Senior Writer and Editor